User Guide - English

FUJITSU

ServerView Suite

ServerView Virtual-IO Manager V3.1

User Guide

Edition October 2012

Comments… Suggestions… Corrections…

The User Documentation Department would like to know your opinion of this manual. Your feedback helps us optimize our documentation to suit your individual needs.

Feel free to send us your comments by e-mail to manuals@ts.fujitsu.com.

Certified documentation according to DIN EN ISO 9001:2008

To ensure a consistently high quality standard and user-friendliness, this documentation was created to meet the regulations of a quality management system which complies with the requirements of the standard DIN EN ISO 9001:2008.

cognitas. Gesellschaft für Technik-Dokumentation mbH

www.cognitas.de

Copyright and trademarks

Delivery subject to availability; right of technical modifications reserved.

All hardware and software names used are trademarks of their respective manufacturers.

Contents

1 Introduction

You use the ServerView Virtual-IO Manager (Virtual-IO Manager or VIOM for short) software to manage the input/output parameters (I/O parameters) of following servers:

- PRIMERGY blade server (BX600, BX400, BX900)

  In Japan, BX600 blade servers are not supported.

- PRIMERGY rack server (RX200 S7, RX300 S7, RX350 S7)
- PRIMERGY tower server (TX300 S7)

When PRIMERGY rack servers are mentioned below, both, the PRIMERGY rack servers and the PRIMERGY tower servers, are meant.

Additionally the LAN connection blade, the Intelligent Blade Panel (IBP) in PRIMERGY blade servers, can be managed via VIOM.

As an extension to the ServerView Operations Manager, it is possible to manage a large number of PRIMERGY blade servers and PRIMERGY rack servers centrally by the central management station using VIOM. This includes virtualizing and, for blade servers, saving the server blade-specific I/O parameters (MAC addresses, WWN addresses, I/O connections including the boot parameters) and configuring and managing a blade server's Intelligent Blade Panel in a hardware-independent server profile.

This server profile can be assigned to a PRIMERGY rack server or server blade:

- For PRIMERGY rack servers: A server profile can be assigned to a PRIMERGY rack server and can also be moved from one PRIMERGY rack server to another.
- For blade servers: The server profile can be assigned to a server blade using VIOM and can also be moved between different server blades of the same or of another blade server.

By assigning the server profiles to a server, you can start the required application without having to reconfigure the SAN and LAN network.

VIOM provides an easy-to-use Web-based graphical user interface, which you can launch using the ServerView Operations Manager. Using this interface, you can carry out all the necessary tasks for managing the I/O parameters of a PRIMERGY blade server or PRIMERGY rack server and for the LAN connection blade, the IBP module in PRIMERGY blade server.

VIOM also provides a comprehensive command line interface, which you can use to perform administrative VIOM tasks in a script-based environment. The VIOM CLI (command line interface) provides an easy-to-use interface for creating scripts and automating administrative tasks.

The command line interface is available both on Windows and Linux platforms, and you install it using separate installation packages. For more information on VIOM CLI, see the documentation entitled "ServerView Virtual-IO Manager Command Line Interface".

1.1 Target groups and objective of this manual

This manual is aimed at system administrators, network administrators and service professionals, who have a sound knowledge of hardware and software. The manual describes the functionality and user interface of theVirtual-IO Manager.

1.2 System requirements

Central management station

- Operating system for the central management station
  - Microsoft Windows® Server™ 2003 all editions
  - Microsoft Windows® Server™ 2003 R2 all editions
  - Microsoft Windows® Server™ 2008 all editions
  - Microsoft Windows® Server™ 2008 R2 all editions
  - Linux Novell (SLES10): SP2 and SP3
  - Novell (SLES 11): SP1 and SP2

- Red Hat RHEL 5.6/5.7/5.8
- Red Hat RHEL 6, 6.1/6.2

In Japan: Novell SLES is not supported.

ServerView Virtual-IO Manager can also be installed in Virtual Machine (VM) under Windows Hyper-V or VMware ESX server. The operating system running on the VM must be one of the above listed operating systems and must be supported by the used hypervisor.

- Installed software packages
  - ServerView Operations Manager as of Version 5.50.13
  - Java Runtime Environment (JRE) version 6.0, update 31 or higher

    Together with ServerView Operations Manager 6.10, it is also possible to use JRE version 7.0, update 7 or higher.
- Fire wall settings
  - Port 3172 must be opened for TCP/IP connection to Remote Connector Service.
  - Port 162 must be opened to receive SNMP traps from iRMC when managing PRIMERGY rack servers.

You can also obtain the current requirements from the release notes. You find the release notes e.g. on a Windows-based management station under **Start - [All] Programs - Fujitsu - ServerView Suite - Virtual-IO Manager - Release Notes**.

License

You must purchase licenses to use the Virtual-IO Manager. At least one license is required. Each license contains a count which determines the allowed number of server profile assigns. If more than one license is registered, the counts are added together.

1.3 Supported Hardware

Managed BX600 blade servers

Supported systems: BX600 S3 with MMB S3. For information on the required firmware version, see the release notes included.

The following table shows which server blades are supported with which range of functions.

Server blade	Scope of functions
BX620 S2, BX620 S3	Server profiles without I/O virtualization but with network connection definition
BX620 S4, BX620 S5, BX620 S6	Server profiles with I/O virtualization and network connection definition
BX630	Server profiles without I/O virtualization but with network connection definition
BX630 S2	Server profiles with I/O virtualization and network connection definition

Table 1: Supported server blades

For information on the BIOS and iRMC firmware version, see the release notes supplied.

The Virtual-IO Manager can only manage BX600 chassis with S3 management blades (MMB S3) that are assembled with the following:

- In fabric 1: IBP or LAN modules
- In fabric 2: IBP modules, LAN modules or FC switch blades of the type SW4016

You must not mix the modules within a fabric.

Fabric 2 can also be empty. Only one of the permitted connection blades can be inserted in fabric 1 and 2 at each time.

In Japan, BX600 blade servers are not supported.

Managed BX400 blade servers

Supported systems: BX400 with MMB S1. For information on the required firmware version, see the release notes supplied.

The following table shows which server blades are supported with which range of functions.

Server blade	Scope of functions
BX920 S2, BX920 S3	Server profiles with I/O virtualization and network connection definition
BX922 S2	Server profiles with I/O virtualization and network connection definition
BX924 S2, BX924 S3	Server profiles with I/O virtualization and network connection definition

Table 2: Supported server blades

For information on the BIOS and iRMC firmware version, see the release notes supplied.

The Virtual-IO Manager can only manage BX400 chassis with S1 management blades (MMB S1) that are assembled with the following:

- In fabric 1:
  - LAN connection blades (PY CB Eth Switch/IBP 1 Gb 36/8+2 (SB11), PY CB Eth Switch/IBP 1 Gb 36/12 (SB11A), PY CB Eth Switch/IBP 1 Gb 18/6 (SB6), or PY CB Eth Switch/IBP 10 Gb 18/8 (SBAX2)) in switch mode or IBP mode, or
  - LAN pass thru connection blades (PY CB Eth Pass Thru 10 Gb 18/18)
- In fabric 2:
  - LAN connection blades (PY CB Eth Switch/IBP 1 Gb 36/8+2, PY CB Eth Switch/IBP 1 Gb 36/12, PY CB Eth Switch/IBP 1 Gb 18/6, or PY CB Eth Switch/IBP 10 Gb 18/8) in switch mode or IBP mode,
  - LAN pass thru connection blades (PY CB Eth Pass Thru 10 Gb 18/18),
  - FC switch blades of the type Brocade 5450, or
  - FC pass thru connection blades (PY CB FC Pass Thru 8 Gb 18/18)
- In fabric 3: same as fabric 2

The LAN connection blades in fabric 3 must run in the same mode. However, only one connection blade can be inserted in fabric 3.

You must not switch the mode of a LAN connection blade if you are using the Virtual-IO Manager to manage the BX400 chassis.

Managed BX900 blade servers

Supported systems: BX900 with MMB S1. For information on the required firmware version, see the release notes supplied.

The following table shows which server blades are supported with which range of functions.

Server blade	Scope of functions
BX920 S1, BX920 S2, BX920 S3	Server profiles with I/O virtualization and network connection definition
BX922 S2	Server profiles with I/O virtualization and network connection definition
BX924 S2, BX924 S3	Server profiles with I/O virtualization and network connection definition
BX960 S1	Server profiles with I/O virtualization and network connection definition

Table 3: Supported server blades

For information on the BIOS and iRMC firmware version, see the release notes supplied.

The Virtual-IO Manager can only manage BX900 chassis with S1 management blades (MMB S1) that are assembled with the following:

- In fabric 1:
  - LAN connection blades (PY CB Eth Switch/IBP 1 Gb 36/8+2 (SB11), PY CB Eth Switch/IBP 1 Gb 36/12 (SB11A), PY CB Eth Switch/IBP 1 Gb 18/6 (SB6), or PY CB Eth Switch/IBP 10 Gb 18/8 (SBAX2)) in switch mode or IBP mode, or
  - LAN pass thru connection blades (PY CB Eth Pass Thru 10 Gb 18/18)
- In fabric 2:
  - LAN connection blades (PY CB Eth Switch/IBP 1Gb 36/8+2, PY CB Eth Switch/IBP 1Gb 36/12, PY CB Eth Switch/IBP 1Gb 18/6, or PY CB Eth Switch/IBP 10 Gb 18/8) in switch mode or IBP mode, or
  - LAN pass thru connection blades (PY CB Eth Pass Thru 10 Gb 18/18),
  - FC switch blades of the type Brocade 5450
  - FC pass thru connection blades (PY CB FC Pass Thru 8 Gb 18/18)
- In fabric 3: same as fabric 2
- In fabric 4: same as fabric 1

The LAN connection blades in a fabric must run in the same mode. However, only one connection blade can be inserted in a fabric.

You must not switch the mode of a LAN connection blade if you are using the Virtual-IO Manager to manage the BX900 chassis.

Managed PRIMERGY rack servers und PRIMERGY tower servers

The following PRIMERGY rack and tower server models are supported:

PRIMERGY model	Scope of functions
RX200 S7	Assign VIOM server profiles with I/O address virtualization and boot configuration for the onboard LAN ports and the supported PCI controller. Network connection definition is not supported.
RX300 S7	
RX350 S7	
TX300 S7	

For information on the BIOS and iRMC firmware version, see the release notes supplied.

The following PCI controllers are supported for all the above PRIMERGY rack server systems:

- Emulex 10GbE OCe10102 CNA
- Emulex 8Gb FC HBA LPe 12002
- Emulex 8Gb FC HBA LPe 1250 (1 channel)
- INTEL 2-port 10GbE (D2755 – Niantec)
- INTEL 4-port (D2745 - Barton Hills)
- INTEL 2-port (D2735 - Kawela 82576NS)
- INTEL Eth Ctrl 4x1Gb Cu PCIe x4 (D3045)
- INTEL Eth Ctrl 2x1Gb Cu PCIe x4 (D3035)
- INTEL 10GbE 10GBase-T(RJ45) PCIe LAN

PCI controller	Scope of functions
Emulex 10GbE OCe10102 CNA	Define physical functions and type (LAN, FCoE, iSCSI) of physical functions. Assign virtual addresses to physical function and optionally define boot parameter. This CNA supports two physical functions for each of both physical ports. The first physical function must be of type LAN. The physical functions for the two physical ports must be defined similarly. This means that the storage function for both physical ports must be of the same type. When using iSCSI with iSCSI boot the iSCSI initiator must be identical. Exception: physical port is completely disabled.
Emulex 8Gb FC HBA LPe 12002 Emulex 8Gb FC HBA LPe 1250 (1 channel)	Assign virtual WWPN and WWNN. Optionally define the first and second boot target and LUN. Disable I/O ports.
INTEL 2-port 10GbE (D2755 – Niantec) INTEL 4-port (D2745 - Barton Hills) INTEL 2-port (D2735 - Kawela 82576NS) INTEL Eth Ctrl 4x1Gb Cu PCIe x4 (D3045) INTEL Eth Ctrl 2x1Gb Cu PCIe x4 (D3035) INTEL 10GbE 10GBase-T(RJ45) PCIe LAN	Assign virtual MAC. Optionally define PXE boot per port. Disable I/O ports. There is no explicit disable functionality for I/O ports in VIOM. I/O ports that are not defined in a VIOM profile will be implicitly disabled if the device supports this functionality.

For information on the required firmware version, see the release notes supplied.

1.4 Changes since the previous edition

The current edition is valid for ServerView Virtual-IO Manager V3.1 and replaces the online manual "PRIMERGY ServerView Suite, ServerView Virtual-IO Manager V3.0", Edition March 2012.

ServerView Virtual-IO Manager V3.1 includes the following new features:

- Support for INTEL Eth Ctrl 4x1Gb Cu PCIe x4 (D3045), INTEL Eth Ctrl 2x1Gb Cu PCIe x4 (D3035), and INTEL 10GbE 10GBase-T(RJ45) PCIe LAN in PRIMERGY rack servers.
- Support of VLAN groups in tagged mode in VIOM server profile (see "Defining networks (LAN) (for blade servers only)" on page 30, "IO-Channels step (Create Server Profile wizard)" on page 190, and "IO-Channels step (Edit Server Profile wizard)" on page 208).
- DCB settings also possible for iSCSI (see " CNA Parameter step (Create Server Profile wizard)" on page 200 and "CNA Parameter step (Edit Server Profile wizard)" on page 218).
- Video Redirection for server blades and rack servers (see "Server Configuration tab" on page 151).
- User-specific display properties are stored session independent (see "Preferences dialog box" on page 238).
- Support of JAVA Runtime Environment 7 (only with ServerView Operation Manager 6.10)
- The section "High Availability - HA" has been updated and now includes VMware HA (see "High-Availability (HA) support" on page 48).

1.5 ServerView Suite link collection

Via the link collection, Fujitsu Technology Solutions provides you with numerous downloads and further information on the ServerView Suite and PRIMERGY servers.

For ServerView Suite, links are offered on the following topics:

- Forum
- Service Desk
- Manuals
- Product information
- Security information
- Software downloads
- Training

The downloads include the following:

- Current software statuses for the ServerView Suite as well as additional Readme files.
- Information files and update sets for system software components (BIOS, firmware, drivers, ServerView agents and ServerView update agents) for updating the PRIMERGY servers via ServerView Update Manager or for locally updating individual servers via ServerView Update Manager Express.
- The current versions of all documentation on the ServerView Suite.

You can retrieve the downloads free of charge from the Fujitsu Technology Solutions Web server.

For PRIMERGY servers, links are offered on the following topics:

- Service Desk
- Manuals
- Product information
- Spare parts catalogue

Access to the link collection

You can reach the link collection of the ServerView Suite in various ways:

1. Via ServerView Operations Manager.

- Select **Help** – **Links** on the start page or on the menu bar.

This opens the start page of the ServerView link collection.

2. Via the ServerView Suite DVD 2 or via the start page of the online documentation for the ServerView Suite on the Fujitsu Technology Solutions manual server.

   You access the start page of the online documentation via the following link:

   http://manuals.ts.fujitsu.com

   - In the selection list on the left, select **Industry standard servers**.
   - Click the menu item **PRIMERGY ServerView Links**.

   This opens the start page of the ServerView link collection.

3. Via the ServerView Suite DVD 1.

   - In the start window of the ServerView Suite DVD 1, select the option **Select ServerView Software Products**.
   - Click **Start**. This takes you to the page with the software products of the ServerView Suite.
   - On the menu bar select **Links**.

   This opens the start page of the ServerView link collection.

1.6 Documentation for the ServerView Suite

The documentation for the ServerView Suite can be found on the ServerView Suite DVD 2 supplied with each server system.

The documentation can also be downloaded free of charge from the Internet. You will find the online documentation at http://manuals.ts.fujitsu.com under the link **Industry standard servers**.

For an overview of the documentation to be found under **ServerView Suite** as well as the filing structure, see the ServerView Suite sitemap (**ServerView Suite** – **Site Overview**).

1.7 Typographic conventions

The following typographic conventions are used:

Convention	Explanation
	Indicates various types of risk, namely health risks, risk of data loss and risk of damage to devices.
	Indicates additional relevant information and tips.
bold	Indicates references to names of interface elements.
`monospace`	Indicates system output and system elements, e.g., file names and paths.
`monospace semibold`	Indicates statements that are to be entered using the keyboard.
blue continuous text	Indicates a link to a related topic.
pink continuous text	Indicates a link to a location you have already visited.
<abc>	Indicates variables which must be replaced with real values.
[abc]	Indicates options that can be specified (syntax).
[key]	Indicates a key on your keyboard. If you need to enter text in uppercase, the Shift key is specified, for example, [SHIFT] + [A] for A. If you need to press two keys at the same time, this is indicated by a plus sign between the two key symbols.

Screenshots

Some of the screenshots are system-dependent, so some of the details shown may differ from your system. There may also be system-specific differences in menu options and commands.

2 Virtual-IO Manager - Introduction

This chapter provides a general introduction to the concept of the Virtual-IO Manager (VIOM).

2.1 Virtual addresses

Physical MAC addresses and WWN addresses are stored on the network card or in the host bus adapter (HBA) of a server blade or PRIMERGY rack server. If a server blade or PRIMERGY rack server has to be exchanged or the operating system and/or the application has to be started on another server, usually the LAN or SAN network has to be reconfigured. This means that whilst the MAC address and the WWN addresses identify a physical server blade, several administrators have to be involved.

To separate the administration areas from each other, it is necessary to keep the I/O parameters (MAC and WWN) outwardly constant.

Using virtual addresses instead of the MAC addresses or WWN addresses stored on the NIC (network interface card) or in the HBA, the addressing remains constant even when a server blade is exchanged at the slot or a PRIMERGY rack server is replaced by another one.

2.2 Special connection blade for blade server

Up to now, blade servers have been used essentially to connect the LAN (Local Area Network) and Fibre Channel ports (FC ports) of individual server blades to the LAN and SAN networks (SAN - Storage Area Network) using switch blades or pass-thru blades, which are inserted in the blade chassis. It is the responsibility of the LAN or SAN administrators to manage these switches. This leads to an overlap of the different administration areas.

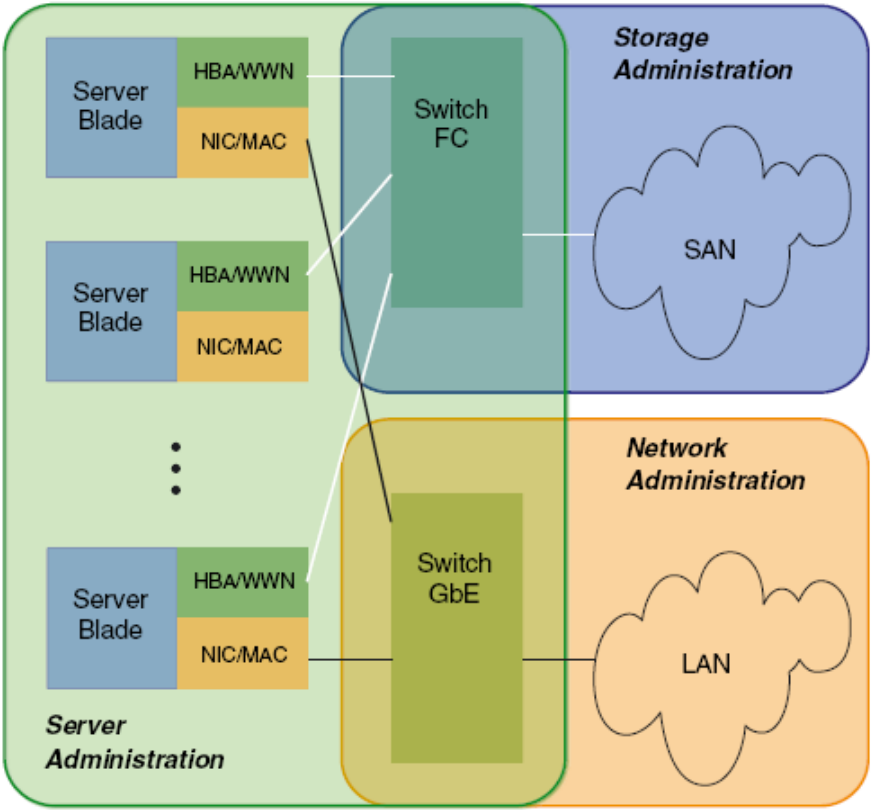

Figure 1: Overlapping areas of responsibility

As the areas of responsibility overlap, this means that up to three administrators may be involved if a server blade's configuration changes, e. g. because a server blade has to be replaced due to hardware problems and, as a result, the switches have to be reconfigured.

The onboard LAN and FC controllers in the server blade are connected to the installed LAN or FC switches via a midplane and are, in turn, connected to the LAN and SAN network via their uplink ports. Providers use specific protocols or protocol extensions for switches from different manufacturers, which can lead to interoperability problems between the internal and external switches of different providers.

To resolve these problems, the switch blades installed in the blade server can be replaced by special connection blades. The following connection blade are available for this:

- For SAN:
  BX600: BX600 4/4Gb FC Switch 12port (SW4016, SW4016-D4) in the Access Gateway mode (FC AG)
  BX400/BX900: Brocade 5450 8 Gb Fibre Channel Switch in the Access Gateway mode (FC AG)
- For LAN:
  BX600: BX600 GbE Intelligent Blade Panel 30/12 or 10/6 (IBP GbE)
  BX400/BX900: Connection Blade PY CB Eth Switch/IBP 1Gb 36/8+2, PY CB Eth Switch/IBP 1Gb 36/12, PY CB Eth Switch/IBP 1Gb 18/6, or PY CB Eth Switch/IBP 10 Gb 18/8 in the IBP mode. (The connection blades can run in switch mode, in IBP mode, or in End Host Mode (EHM).)

These connection blades offer the advantage of a switch (cable consolidation) without the above-mentioned disadvantages.

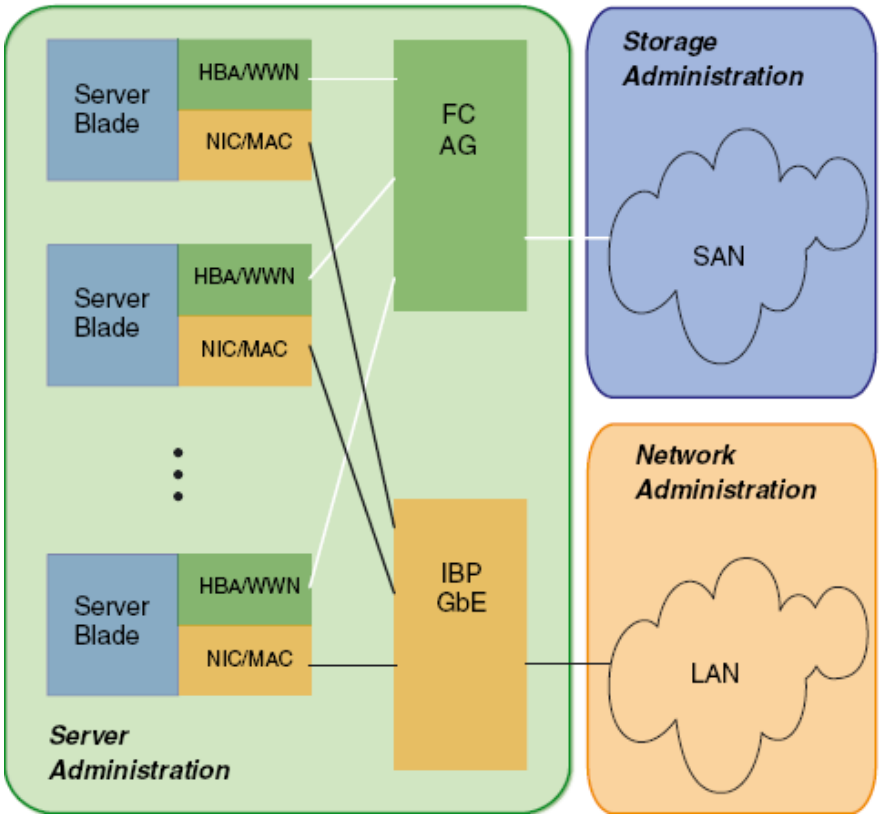

Figure 2: Separate areas of responsibility

2.3 Management with VIOM - Procedure

You use the ServerView Virtual-IO Manager (VIOM) to manage the connection blades of a blade server and to maintain the relevant I/O parameters constant at the chassis slot of a blade server or at the PRIMERGY rack server. VIOM is installed on the central management station and integrated in the ServerView Operations Manager.

For blade servers, management with VIOM essentially includes the following functions:

- Defining the network paths on the Intelligent Blade Panel (IBP module)
- Defining the I/O parameters (including virtual addresses) of a server blade
- Saving the I/O parameters in combination with the required network paths in server profiles
- Assigning the server profiles to the server blades or to empty slots as well as moving a server profile between any number of server blades

You can assign the server profiles to any number of different server blades of a blade server or even to another blade server, provided that the required network connections are available on the respective chassis.

For PRIMERGY rack servers, management with VIOM essentially includes the following functions:

- Defining the I/O parameters (including virtual addresses) of a PRIMERGY rack server
- Saving the I/O parameters in server profiles
- Assigning the server profiles to a PRIMERGY rack server

Before you can execute the functions above, a blade server chassis or PRIMERGY rack server must be managed by VIOM. You also do this using the GUI of the Virtual-IO Manager (see chapter "Managing servers with VIOM" on page 257).

Management using VIOM is divided into the following key steps:

1. Before VIOM can work with a blade server chassis or PRIMERGY rack server it must be managed by VIOM.
2. For blade servers with IBP, you can define external network connections.
3. You then define the corresponding profiles for all applications/images and save them in the server profile repository on the central management station.

4. You can then assign these server profiles to any of the individual slots of a blade server or to a PRIMERGY rack server.
5. If required, you can remove the assignment of the server profile.

   For blade server, you can move the profiles from one blade server slot to another, or move them to another blade server. For PRIMERGY rack servers, you can move the profiles from one server to another server.

2.4 Defining networks (LAN) (for blade servers only)

In order for VIOM to be able to switch the network paths correctly when assigning profiles, first VIOM has to know which networks are present on the respective chassis on which uplink ports.

This also makes it possible to separate individual server blades or groups of server blades from a network perspective so that two server blade groups do not have any connection to each other in terms of network.

Defining network paths on an IBP module includes the following steps:

- Defining an uplink set.

  An uplink set comprises one or several uplink ports. An uplink port is an external port that connects the chassis with your LAN infrastructure. If an uplink set is used by several virtual network connections (VLANs), the uplink set is referred to as a shared uplink set.
- Defining one or several networks that are assigned to the uplink set.

  

  The definition of a network in the context of VIOM, refers to the allocation of a meaningful name for network access from outside the network

If a blade server chassis is managed by VIOM, manual configurations (not done by VIOM) of an IBP connection blade are not supported. Manual configuration of IBP connection blades might result in incorrect behavior of VIOM or get lost during configuration by VIOM. Before managing a chassis by VIOM IBP connection blades should be set to factory default setting except IP address configuration of the administrative interface, user assigned name and access protocol (SSH or telnet).

By default, the IBP module (IBP 10/6) is supplied with the following configuration.

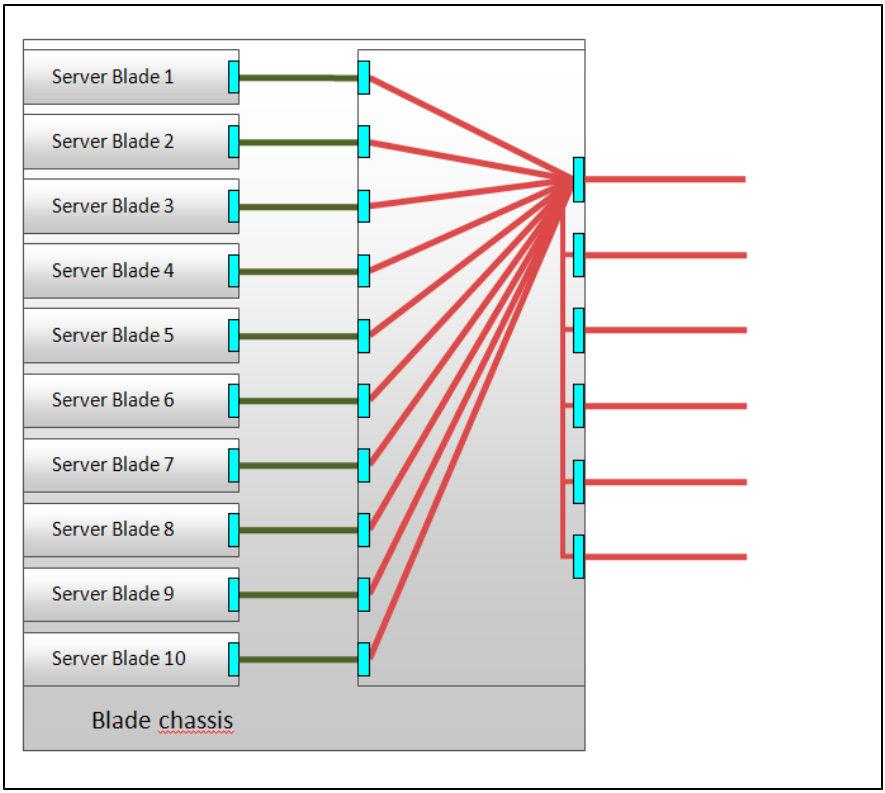

Figure 3: Standard configuration of the IBP module (10/6)

All uplink ports of the IBP module (IBP 10/6) are combined in one uplink set.

In the case of IBP 30/12, the first 8 uplink ports are combined in one uplink set by default, and all 30 downlinks are connected with this standard uplink set.

Using VIOM, you can change the standard configuration of an IBP module. You can combine several uplink ports into one uplink set as well as define several uplink sets for a LAN connection blade. This gives you several independent network paths e. g. for different applications (e. g. database server, communication server) or individual areas (e. g. development, accounting or personnel administration).

To find out what happens when you activate the management of a blade server using the Virtual-IO Manager with the standard IBP configuration, see section"VIOM internal operations on blade servers" on page 261 .

The following figure provides an overview of typical uplink sets that you can configure using VIOM.

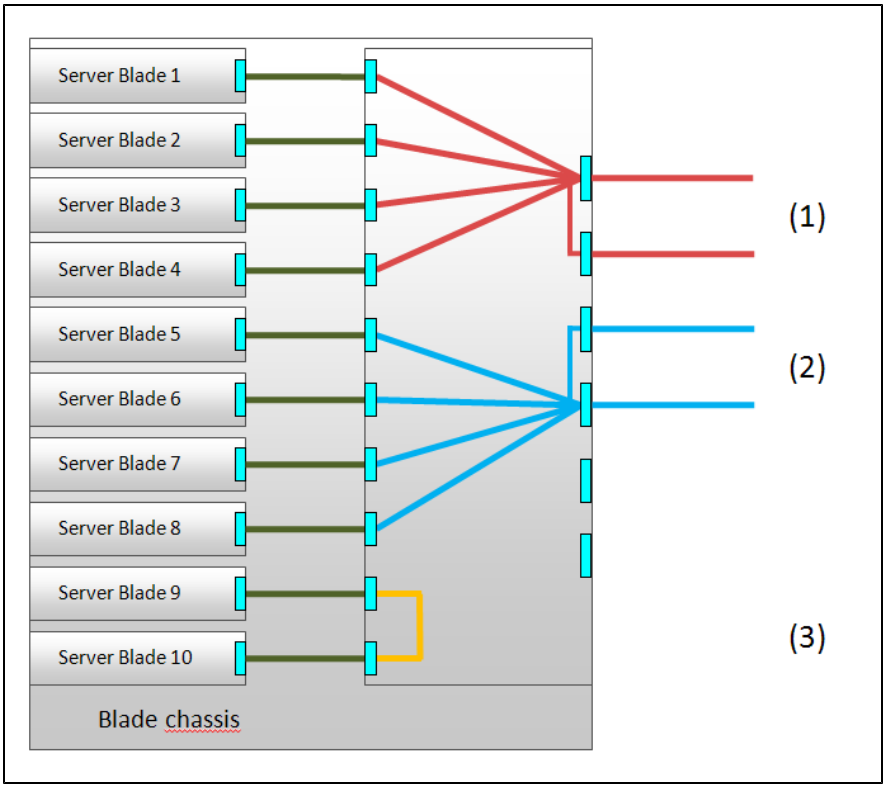

Figure 4: Typical uplink sets

The uplink ports can be assigned to an uplink set as active ports or as back-up ports. As a result, there are different ways of configuring an uplink set:

- "Port backup" configuration
  When you configure a "port backup", you define an uplink set with at least two uplink ports, and configure one of these as an active port and the other as a backup port. In this case, the active port switches to the backup port if an error occurs (linkdown event for all active ports). In the figure shown, this could be the uplink sets (1) and (2) if for each of these one port of the uplink set is configured as an active port and one port as a backup port.
- Link aggregation group
  By grouping several active uplink ports in one uplink set, a link

aggregation group (LAG) is formed. By providing several parallel connections, you achieve higher level of availability and a greater connection capacity. In the figure above, this could be the uplink sets (1) and (2) if both uplink ports of these uplink sets are configured as active ports. If an uplink set has several backup ports, these backup ports also form a link aggregation group automatically in the case of a failover. It is essential that the ports of an external LAN switch, which are linked to a LAG, form a static LAG.

Using VIOM, it is possible to define a number of networks:

- Internal networks (**Internal network**)
- Single networks (**Single network**)
- Virtual networks with VLAN IDs (**VLAN networks**)
- Virtual networks with VLAN IDs/native VLAN ID **(VLAN networks**)
- Service LAN (**Dedicated service network**)
- Service VLAN (**Service VLAN networks**)

The network types in bold indicate their corresponding names in the VIOM GUI.

Internal networks

An internal network refers to a network connection within the IBP, in which server blades are only linked to each other. However, no uplink ports are assigned to this network connection.

In this case, it is an internal connection via the IBP module.

It makes sense to have an internal network if the server blades only need to communicate amongst each other and, for security reasons, there must be no connection to an external network.

In the figure, (3) represents an internal network.

"Single" networks

VIOM interprets a "single" network as an uplink set that is only used for access in one network.

A key attribute of a "single" network is that it is VLAN transparent. You can therefore channel several external networks with different VLAN tags (or also without VLAN tags) through a "single" network.

Packets with or even without a VLAN tag, which arrive at the uplink ports from outside the network, are channeled to the related server blades with the corresponding network. The same applies to the network packets that come from the server blades.

In the figure, (1) and (2) illustrate "single" networks.

Virtual networks with VLAN IDs

Depending on the IBP module, you have 6 to 12 uplink ports available. You can define as many different networks as there are uplink ports. If networks are to be created with backup ports or with a link aggregation group, then the number of possible networks on an IBP module is automatically reduced. You can get around this restriction regarding the uplinks that are physically available by defining virtual networks (Virtual Local Area Network - VLAN).

By setting up virtual networks, which can be identified by unique numbers known as VLAN IDs, you can set up several logical networks that are completely separate from each other from a technical and network perspective. These networks share an uplink set ("shared uplink set") without the server blades of one virtual network being able to communicate with server blades of the other virtual networks.

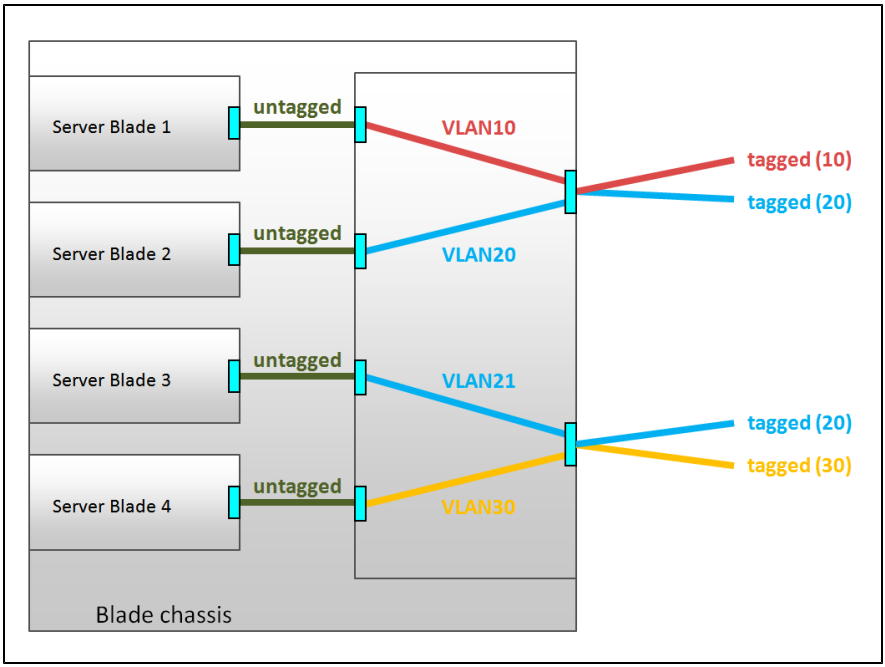

Figure 5: Networks with VLAN ID

In the figure above, two shared uplink sets are configured on the IBP module. Two virtual networks **VLAN10** and **VLAN20** with the VLAN IDs 10 and 20 are assigned to the upper shared uplink set, and two virtual networks **VLAN21** and **VLAN30** with the VLAN IDs 20 and 30 are assigned to the lower shared uplink set. Although both uplink sets have virtual networks with the VLAN ID 20, these are two different virtual networks. Server blade 2 cannot communicate with server blade 3.

Packets that come from outside the network, which have a VLAN tag that corresponds to the VLAN ID of a virtual network, are transferred in precisely this VLAN network. Before the packet exits the module on the server blade side, the VLAN tag of a virtual network is removed in the same way as a "port-based" VLAN.

Packets that come from outside the network, which have a VLAN tag that does not match any VLAN ID of a virtual network, are not transferred. They are dropped.

Packets that come from outside the network with no VLAN tag are also dropped. This behavior can be changed by configuring a virtual network as native VLAN (see "Virtual network with a VLAN ID as native VLAN" on page 37).

Packets that come from a server blade, which do not have a VLAN tag, are routed in the VLAN network to which the LAN port of the server blade is connected. In the process, VLAN tags with the VLAN ID of the virtual network are added to these packets. These packets exit the IBP module at the uplink ports of the related uplink set with this VLAN tag.

Packets that come from a server blade, which have a VLAN tag, are not transferred to a VLAN network. They are either dropped or transferred elsewhere (e.g. to service networks). But VLAN networks can also be used with tagged packets when the network is used in tagged mode (see "Virtual networks with VLAN ID used in tagged mode" on page 38).

Virtual network with a VLAN ID as native VLAN

You can select a virtual network of a shared up link set as the default or "native" VLAN. All packages that do not contain a VLAN ID will be allowed through this connection.

Packets that come from outside the network, which do not have a VLAN tag, are routed in the network with the native VLAN ID and assigned a corresponding VLAN tag in the process.

Packets that come from outside the network, which have a VLAN tag that corresponds to the native VLAN ID, are not transferred in any of the networks belonging to the uplink set. They are dropped.

Packets that come from a server blade to the native VLAN network, exit an IBP module without a VLAN tag. The VLAN tag is therefore removed from the network packet before it exits the IBP module at the uplink port.

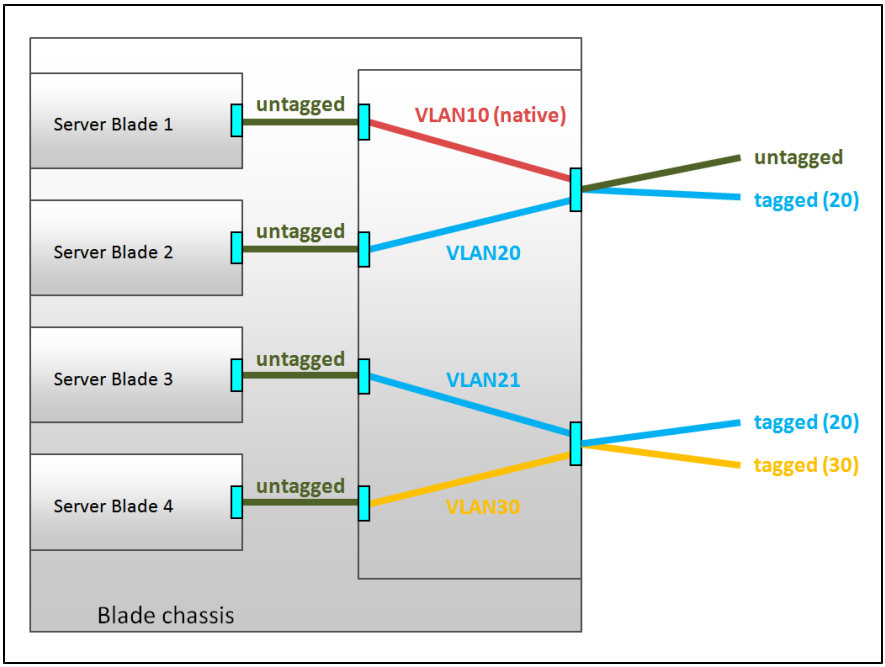

Figure 6: Networks with a VLAN ID and a native VLAN ID

In the figure above, the VLAN ID 10 is defined as the native VLAN ID in the upper shared uplink set. As a result, the data packets of server blade 1 with the VLAN ID 10 (red) exit the uplink without a VLAN ID tag. Incoming data packets without a VLAN ID tag are assigned the VLAN ID 10 internally. These data packets are only transferred to server blade 1.

Virtual networks with VLAN ID used in tagged mode

While normally the VLAN IDs of packets that leave the IBP on the server blade side are removed, it is possible to use a VLAN network in tagged mode. This means that all packets retain their VLAN tag when they are transmitted to the server blade. Packets that arrive on the server blade side of the IBP must have a VLAN tag with the corresponding VLAN ID if they are to transferred to this VLAN network in tagged mode.

Packets without a VLAN tag are dropped unless there is a VLAN network in untagged mode associated with the same downlink.

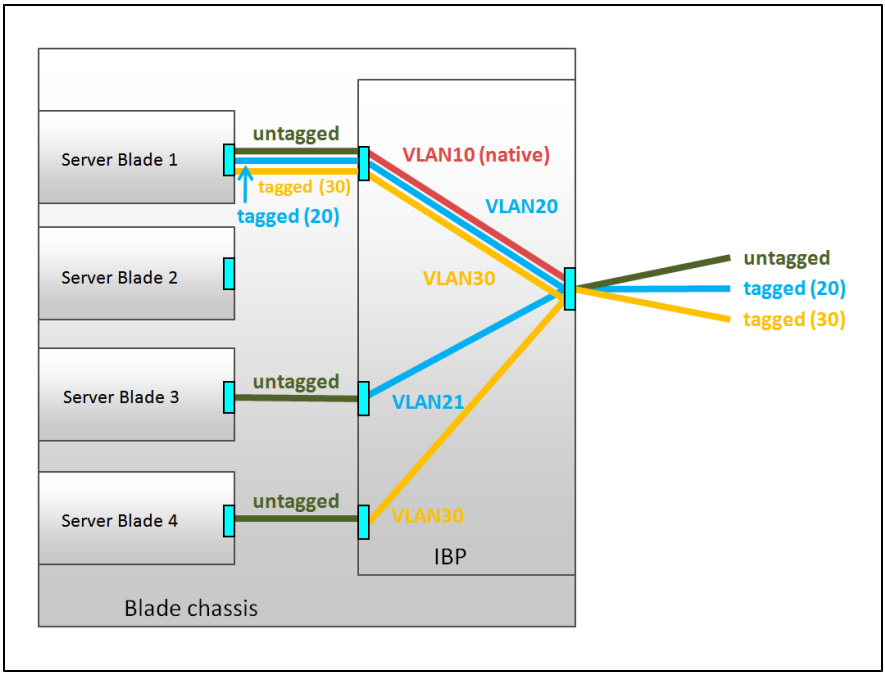

Figure 7: Virtual networks with VLAN ID used in tagged mode

Several VLAN networks in tagged mode can be used on the same IBP downlink port. They can also be combined with service networks.

The mode in which a VLAN network is used is controlled by network definitions in a server profile. It cannot be specified within the network settings.

The advantage of the tagged mode is that the same VLAN networks can be used either untagged for separate server blades on different downlinks or tagged for one server blade with separate virtual machines on one downlink.

Dedicated service networks

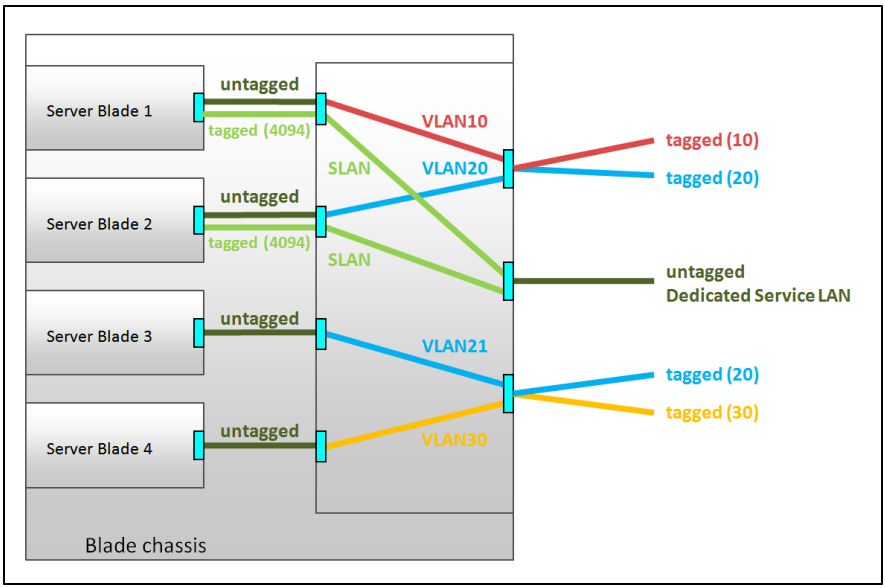

Figure 8: Dedicated service networks

The dedicated service network is designed to separate LAN traffic of an iRMC from the operating system LAN traffic if the iRMC is not using a separate management LAN but is configured to share its LAN traffic with an onboard LAN port of the server blade. In order to separate the LAN traffic of iRMC and operating system in this case, the iRMC must also be configured to use a VLAN tag for its LAN packets. A dedicated service network defined with the same VLAN ID as used by the iRMC allows the tagged iRMC LAN packets to be routed to specific uplink port(s) (external port(s)), whereas the other LAN packets from the operating system are routed to a separate uplink port.

In addition, the dedicated service network can also be used to route the LAN packages of a virtual NIC defined in the operating system running on the server blade to specific uplink ports. In order to do this, the virtual NIC in the operating system must be configured to send all packets with a VLAN tag.

The same VLAN tag must be specified when defining the dedicated service network that is to transport these packages.

The behavior of a dedicated service network is such that it receives tagged packets from the server blade, but the tags are stripped when they leave the uplink port. Incoming untagged packets at the uplink port are tagged and sent to the corresponding downlink ports (internal ports)/blade server as tagged packets. Incoming tagged packets at uplink ports are dropped.

Note that dedicated service networks may overlap on the downlink ports with single networks, VLAN networks, other dedicated service networks, and Service VLAN networks (explained below). The untagged packets received from the server blade or uplink port should obey the rule of the single network or VLAN network that overlaps with the dedicated service network.

The VLAN tags of the overlapping VLAN networks, dedicated service networks and Service VLAN networks must be different.

Dedicated service networks cannot overlap with any other network at the uplink ports. This means the uplink ports of a dedicated service network can only be assigned to this dedicated network.

Service VLAN networks

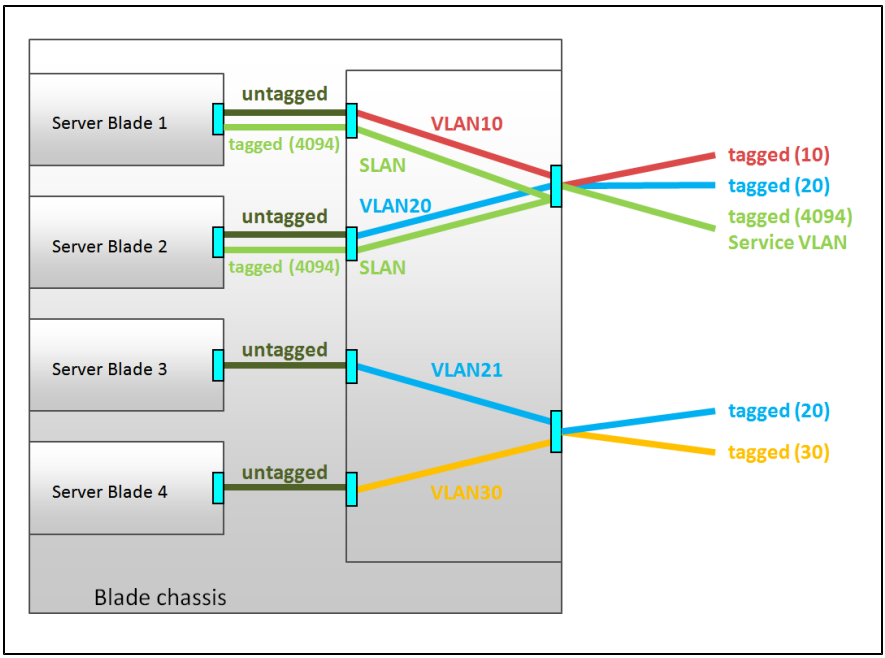

Figure 9: Service VLAN networks

The Service VLAN networks are designed to separate LAN packages of multiple virtual NICs defined in the operating system running on the server blade and route them to specific uplink (external) ports. To do this, the different virtual NICs in the operating system must be configured to send their packages with a VLAN tag that is identical to the Service VLAN ID of the Service VLAN network.

The behavior of a Service VLAN network is such that it receives tagged packets with the Service VLAN ID from the server blade and forwards them to uplink ports as tagged packets. The LAN packages leave the IBP tagged at the uplink ports. Incoming tagged packets with the Service VLAN ID (at the uplink port) are sent to the corresponding downlink (internal) ports/blade servers as tagged packets.

Note that Service VLAN networks may overlap on the downlink ports (with single networks, VLAN networks, dedicated service networks and other Service VLAN networks). The untagged packets received from the server blade or uplink port should obey the rule of the overlapping single network or VLAN network.

The VLAN tags of the overlapping VLAN networks, dedicated service networks and Service VLAN networks must be different.

Different Service VLAN networks may share the same uplink ports. If the port that is member of the Service VLAN network receives tagged packets with the Service VLAN ID (SVID) of a specific Service VLAN network, these received tagged packets will be forwarded based on the definition of this Service VLAN network. The Service VLAN networks with disjoint uplink sets may have identical SVIDs.

Service VLAN networks may also share the same uplink ports with VLAN networks. The VLAN tag of Service VLAN networks and VLAN networks sharing the same uplink ports must be different.

2.5 Server profiles

A server profile contains the following VIOM-specific parameters:

1. Defining the connection in external networks (see section "Defining network paths (LAN)" on page 277), only for blade servers
2. Defining the physical identity in the form of I/O addresses (MAC, WWN)
3. Defining the boot devices with parameters

To activate a server profile of this type, it must be assigned to a server blade slot or a PRIMERGY rack server.

For a blade server, it can be moved to another slot if required (e. g. in the event of server blade failure). The server blade in another slot thereby assumes the identity of the previous server blade. In this way, server profiles allow the available blade hardware to be used flexibly.

In this context, VIOM provides the option to define a slot as a spare slot. If a problem occurs or maintenance work needs to be carried out, you can trigger

a server profile failover, which searches for a suitable spare server blade that will assume the tasks of the failed server blade.

To use server profiles, you must do the following in the Virtual-IO Manager:

1. Define a server profile
2. Assign the profile to a slot or a PRIMERGY rack server

2.5.1 Defining server profiles

A server profile is made up of a set of parameters that contain the related VIOM parameters. These include:

- (Virtual) MAC addresses and WWN addresses
- Boot parameters
- For blade servers only: LAN connections for the I/O channels of a server blade

The server profiles are stored centrally and independently of hardware under a user-defined name in a server profile repository on the central management station.

2.5.2 Assigning server profiles

The server profiles that are stored in the central server profile repository can be assigned to the slots of a blade server or to a PRIMERGY rack server using VIOM. In order to do this, the blade server or PRIMERGY rack server must be managed by VIOM. In addition, you must switch off the server blade in the corresponding slot or the PRIMERGY rack server in order to assign a server profile to this slot or PRIMERGY rack server.

For blade servers, a server profile can also be assigned to an empty slot. A slot can thus be prepared for use at a later date. Using virtual addressing, you can, e. g. quickly replace a faulty server blade by preconfiguring another server blade without changing the configuration.

2.5.3 Dedicated LAN connections (only for blade servers)

You can assign each I/O port of a server blade to an explicit network in the server profile. As a server profile is not connected to any hardware, only the network name is recorded in it.

If a server profile is assigned to a slot, the downlinks connected to the I/O channels of the slot are added to the IBP modules in the specified network. The networks explicitly named in the server profile must be configured beforehand in the affected IBP modules.

If non-VIOM capable LAN modules are installed (Open Fabric mode), you cannot set any dedicated LAN connections (paths). In this case, you must work with profiles whose I/O ports do not contain any network assignment.

2.5.4 Virtualizing I/O parameters

Virtualizing the physical server identity in the form of physical MAC addresses, WWN addresses and boot parameters is a key function of the ServerView Virtual-IO Manager software.

By defining virtual I/O addresses and boot parameters as part of a server profile, you can easily move an operating system image or an application from one server blade or PRIMERGY rack server to another.

The following basic I/O parameters belong to the virtualization parameters:

- Virtual MAC address (LAN)
- Virtual WWN addresses (Fibre Channel)

You can also define the iSCSI boot parameters for LAN ports which are defined as iSCSI boot devices. For each Fibre Channel HBA port the following SAN boot configuration parameters can be virtualized:

- Boot
- 1st target port name (WWPN of the target device)
- 1st target LUN
- 2nd target port name (WWPN of the target device)
- 2nd target LUN

Blade Servers

The virtualization I/O parameters of all the server blades of a chassis are stored in a specific table in the management blade (MMB) of this blade server. When a server blade powered on, checks are run in the boot phase to determine whether virtualization parameters are defined in the MMB table for this server blade slot. These parameters are transferred to the I/O adapters so that the virtualized addresses are used in the same way as the physical addresses assigned by the manufacturer. This ensures that no changes need to be made if a server blade is exchanged or a server profile moved.

If a server blade or a mezzanine card is removed from a blade server and inserted in the slot of another blade server that is not managed by VIOM, then the physical I/O addresses assigned by the manufacturer will be used automatically. The same applies if the virtualization of the I/O addresses for a slot is switched off, e. g. if the corresponding server profile is moved.

If the central management server is switched off or the connection between the management station and the management blade is interrupted, all the blade servers use the configuration last defined.

Once the connection to the external networks is configured and the server profiles assigned with virtualization parameters by the ServerView Virtual-IO Manager, the management station does not necessarily have to run with the Virtual-IO Manager software. To operate the "virtualized" blade server chassis, the software is not required.

PRIMERGY rack servers

The virtualization I/O parameters of a PRIMERGY rack server are stored in a specific table in the baseboard management controller (iRMC) of the server. When a PRIMERGY rack server is powered on, checks are running in the boot phase to determine whether virtualization parameters are defined in the iRMC table. These parameters are transferred to the I/O adapters so that virtualized addresses are used in the same way as the physical addresses assigned by the manufacturer.

If the virtualization of the I/O addresses for a slot is switched off, e. g. if the corresponding server profile is unassigned, the physical I/O addresses assigned by the manufacturer will automatically be reactivated in the next boot phase.

The iRMC of a PRIMERGY rack server loses virtualization I/O parameter table during power failures. So the table has to be rewritten by ServerView Virtual-IO Manager before the server is powered on again. This restoration process is done automatically. But this requires that the management station has to be kept running as long as PRIMERGY rack servers are managed. For further information, see "VIOM-internal operations on a PRIMERGY rack server" on page 266.

2.6 Server profile failover (for blade servers only)

If a problem occurs or maintenance work needs to be carried out, VIOM provides the option to move the server profiles from the affected server blade to a suitable server blade within the same blade server.

In order to do this, you must define spare slots that assume the tasks of the other server blade in such a case. It is advisable to install server blades at the spare slots so that they are available if a problem occurs or maintenance work needs to be carried out. A failover of this type can only take place if the server blade on which the failover is to take place is switched off.

If a server blade fails, for example, you launch the failover function via the context menu of the corresponding server blade. VIOM then searches for a spare slot that has a server blade to which the server profile can be assigned. Once such a slot has been found, the profile assignment on the affected server blade is deleted, and the server profile is assigned to the new server blade. The new server blade thus assumes the role of the failed server blade including the network addresses.

The Virtual-IO Manager does not make any changes to the boot image in a SAN and does not clone any disk images to the local hard disk of the replacement server blade.

2.7 High-Availability (HA) support

VIOM supports the following high-availability environment:

- Windows 2008 R2 Hyper-V cluster with ServerView Operations Manager and ServerView Virtual-IO Manager installed on a virtual machine with Windows Server operating system.
- VMware HA with ServerView Operations Manager and ServerView Virtual-IO Manager installed on a virtual machine with Windows Server operating system.

This means that the ServerView management station is a virtual machine running on a Windows 2008 Hyper-V cluster or in a VMware HA environment.

High availability of Hyper-V cluster

The following Hyper-V high-availability configurations will be supported:

Operating system	**Admin server if HA**	
	Guest OS	**Hypervisor**
Windows Server 2003 R2 Enterprise (x86, x64) SP2 or higher	✓	---
Windows Server 2003 R2 Standard (x86, x64) SP2 or higher	✓	---
Windows Server 2008 R2 Datacenter [*]	✓	✓ [Hyper-V]
Windows Server 2008 R2 Enterprise [*]	✓	✓ [Hyper-V]
Windows Server 2008 R2 Standard [*]	✓	✓ [Hyper-V]

Operating system	Admin server if HA	
	Guest OS	Hypervisor
Windows Server 2008 R2 Foundation [*]	✓	✓ [Hyper-V]
Windows Server 2008 Standard (x86, x64) [*]	✓	✓ [Hyper-V] (only x64)
Windows Server 2008 Enterprise (x86, x64) [*]	✓	✓ [Hyper-V] (only x64)

Figure 10: Supported Hyper-V high-availability configurations

[*] The Windows Server Core Installation option is not supported for admin server and guest OS on VM.

To set up the Windows 2008 Hyper-V cluster and the virtual machine that will be controlled from it, click here for the Microsoft instructions:

http://technet.microsoft.com/en-us/library/cc732181%28v=ws.10%29.aspx

If there is a fault in the Hyper-V cluster node, the Microsoft cluster will perform a failover action of the Hyper-V environment to the other cluster node and restart the virtual machine that is acting as the ServerView Suite management station.

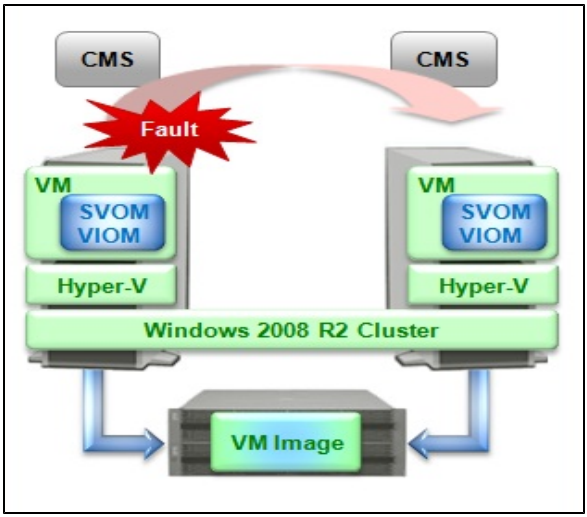

Figure 11: Failover action of the Hyper-V environment to the other cluster node

In the failover clustering of the Hyper-V environment, ServerView supports the cold migration of virtual machines.

To setup the Hyper-V cluster, proceed as follows:

On the **primary** node:

1. Connect with shared storage.
2. Configure BIOS.
3. Install Hyper-V roles.
4. Install and configure EMC Solutions Enabler (if used).
5. Add a failover clustering function.
6. Create a Hyper-V virtual network.
7. Create clusters.
8. Prepare virtual machines.
9. Register virtual machines in clusters.
10. Install and configure storage management software.
11. Install and configure VM management software.

12. Install and configure ServerView Operations Manager and ServerView Virtual-IO Manager.

On the **secondary** node:

1. Connect with shared storage.
2. Configure BIOS.
3. Install Hyper-V roles.
4. Install and configure EMC Solutions Enabler (if used).
5. Add a failover clustering function.
6. Create a Hyper-V virtual network.
7. Install Hyper-V roles.
8. Add a failover clustering function.
9. Create a Hyper-V virtual network.
10. Create clusters.
11. Prepare virtual machines.
12. Register virtual machines in clusters.
13. Operate the management station in a cluster.

For details of items 7 to 13, refer to the Hyper-V manual.

If an error occurs on a VM guest, the operation will continue if the VM guest is switched over.

High availability of VMware HA

To make use of the high-availability functionality of VMware HA, you must use the operating system VMware Infrastructure 3 with the two concepts Cluster and Resource Pool.

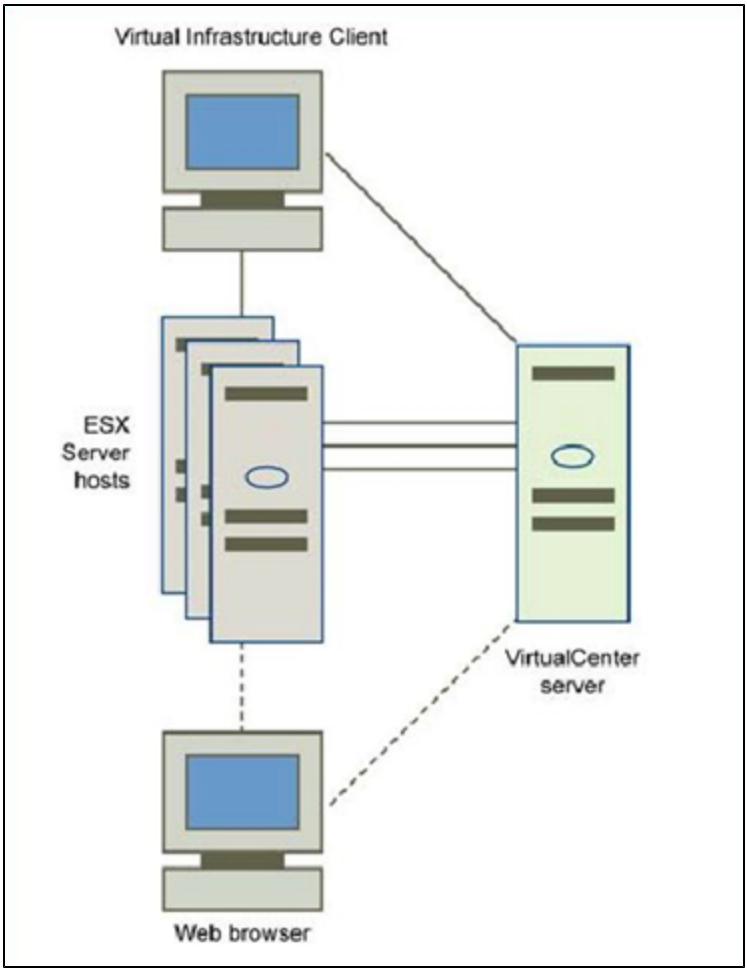

Figure 12: Architecture and typical configuration of VMware Infrastructure 3

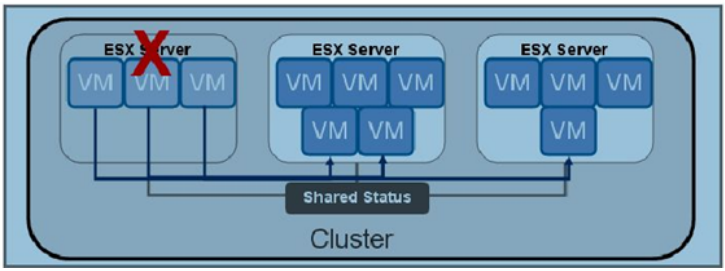

Figure 13: Host failover with VMware HA

VMware HA links up multiple ESX/ESXi servers to form a cluster with shared resources. If one host fails, VMware HA reacts immediately by restarting any affected virtual machine on a different host. The cluster is created and managed via VirtualCenter.

For a detailed description of the high-availability functionality with VMware HA, visit http://www.vmware.com/pdf/vmware_ha_wp.pdf.

HA functionality supported by Virtual-IO Manager

HA functionality is supported by Virtual-IO Manager by the following measures:

- The central ServerView Virtual-IO Manager service is configured to be restarted automatically in the event of failure. This automatic restart of the service is configured during installation of Virtual-IO manager. By default the following restart behavior is configured:
  - The first restart of the service is tried 5 seconds after unexpected termination of the service.
  - The second restart is tried 30 seconds after termination of the service.
  - Subsequent restarts are tried 60 seconds after termination of the service.
  - The restart counter is reset after 600 seconds.

The details of this configuration cannot be seen via the normal graphical user interface of the service manager. To see the details, you must use the command line interface of the service manager. The command `sc qfailure ServerViewVirtualIOManagerService` displays the current configuration, where **ServerViewVirtualIOManagerService** is the name of the service.

- The ServerView Virtual-IO Manager Backup service is also configured for automatic restart in the event of failure during installation.

  

  By default the Virtual-IO Manager Backup service is not configured to start automatically, as it requires configuration.

  By default the following restart behavior is configured:

  - The first restart of the service is tried 5 seconds after unexpected termination of the service.
  - The second restart is tried 30 seconds after termination of the service.
  - Subsequent restarts are tried 120 seconds after termination of the service.
  - The restart counter is reset after 600 seconds.

  

  The details of this configuration cannot be seen via the normal graphical user interface of the service manager. To see the details, you must use the command line interface of the service manager. The command `sc qfailure ServerViewVirtualIOBackupService` displays the current configuration where **ServerViewVirtualIOManagerService** is the name of the service.

- If the Virtual-IO Manager is interrupted during a configuration request (for example, creation of networks in an IBP connection blade or assignment of a VIOM server profile) while executing configuration commands on hardware modules, Virtual-IO Manager will undo the changes already made the next time the service starts. This means that, when the serv-

ice has restarted, the configuration should be the same as it was just before the interrupted request.

Some configuration actions of the Virtual-IO Manager user interface consist of several “independent” internal configuration requests. The Virtual-IO Manager can only undo the last internal configuration request.

If Virtual-IO Manager successfully executes all necessary changes for a request but is just interrupted while sending the response to the Virtual-IO Manager client, the changes will not be undone.

- The transaction concept of the Virtual-IO Manager should allow you to restart the service and execute the described undo actions if the database is not corrupted by the SQL database service used.

What Virtual-IO Manager does not do

Virtual-IO Manager does not control the availability of the ServerView Virtual-IO Manager service. It also does not check the availability of the virtual machine that is used as the ServerView Suite management station. The letter should be done by the Microsoft Hyper-V cluster if it is correctly configured.

3 Installation and uninstallation

You can install the Virtual-IO Manager on a central management station (CMS) under Windows or Linux (see section "Installing the Virtual-IO Manager on a Windows-based CMS" on page 58 and "Installing the Virtual-IO Manager on a Linux-based CMS" on page 73).

Please check first the requirements for installing the Virtual-IO Manager on CMS (see section "Prerequisites for the VIOM installation" on page 57).

If a previous version is already installed on the management station, an update installation runs automatically when you install the new version. All previous VIOM configurations and definitions remain the same (see section "Updating the Virtual-IO Manager on a Windows-based CMS" on page 72) and "Updating the Virtual-IO Manager on a Linux-based CMS" on page 89.

If you want to use the command line interface of the Virtual-IO Manager (VIOM CLI), you must install the VIOM CLI software package. You will find details on how to install and use VIOM CLI in the "Virtual-IO Manager Command Line Interface" manual.

3.1 Prerequisites for the VIOM installation

The requirements for installing the Virtual-IO Manager on a central management station are as follows:

- Operating system for the central management station
  - Microsoft Windows® Server™ 2003 all editions
  - Microsoft Windows® Server™ 2003 R2 all editions
  - Microsoft Windows® Server™ 2008 all editions
  - Microsoft Windows® Server™ 2008 R2 all editions
  - Linux Novell (SLES10): SP2 and SP3
  - Novell (SLES 11): SP1 and SP2
  - Red Hat RHEL 5.6/5.7/5.8
  - Red Hat RHEL 6, 6.1/6.2

In Japan: Novell SLES is not supported.

ServerView Virtual-IO Manager can also be installed in Virtual Machine (VM) under Windows Hyper-V or VMware ESX server. The operating system running on the VM must be one of the above listed operating systems and must be supported by the used hypervisor.

- Installed software packages
  - ServerView Operations Manager as of Version 5.50.13
  - Java Runtime Environment (JRE) version 6.0, update 31 or higher

    Together with ServerView Operations Manager 6.10, it is also possible to use JRE version 7.0, update 7 or higher.
- Fire wall settings
  - Port 3172 must be opened for TCP/IP connection to Remote Connector Service.
  - Port 162 must be opened to receive SNMP traps from iRMC when managing PRIMERGY rack servers.

You can also obtain the current requirements from the release notes. You find the release notes e.g. on a Windows-based management station under **Start - [All] Programs - Fujitsu - ServerView Suite - Virtual-IO Manager - Release Notes**.

3.2 Installing the Virtual-IO Manager on a Windows-based CMS

The corresponding software is supplied with the PRIMERGY ServerView Suite DVD1. You can find the entire software for the PRIMERGY ServerView Suite under **ServerView Software Product Selection**. To find the Virtual-IO Manager software package **SV_VIOM.exe** in this product selection, choose **ServerView – Virtual-IO Manager**.

3.2.1 Installing the Virtual-IO Manager using a graphical interface

Installation process

1. Insert the PRIMERGY ServerView Suite DVD 1 in the DVD-ROM drive. If the DVD does not start automatically, click the **setup.exe** file in the root directory of the DVD-ROM.
2. Select the option **ServerView Software Products**.
3. Click **Start**.
4. In the next window, select the required language.
5. Select **ServerView – Virtual-IO Manager**.
6. Double-click the **SV_VIOM.exe**. The installation wizard is launched. After determining a number of parameters of the existing operating system base, the following window is displayed:

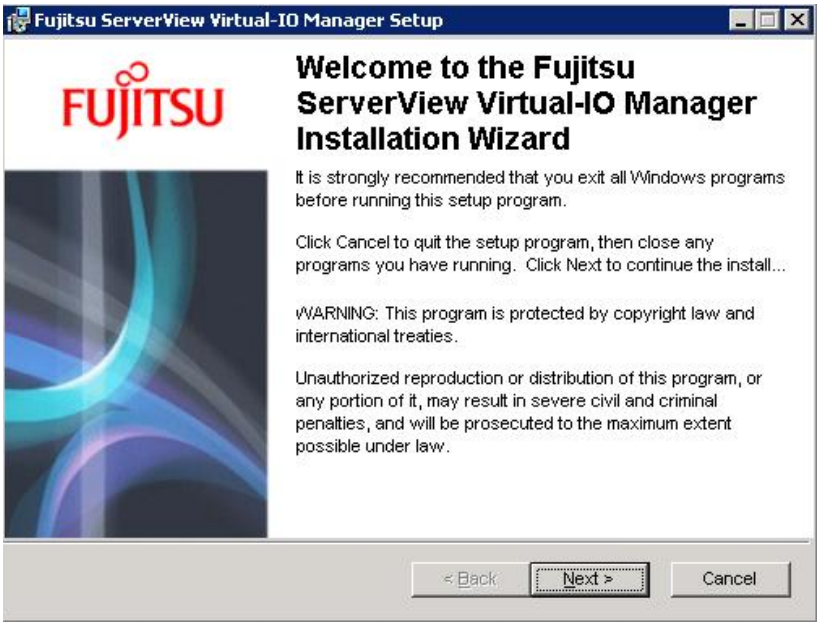

7. Click **Next**.

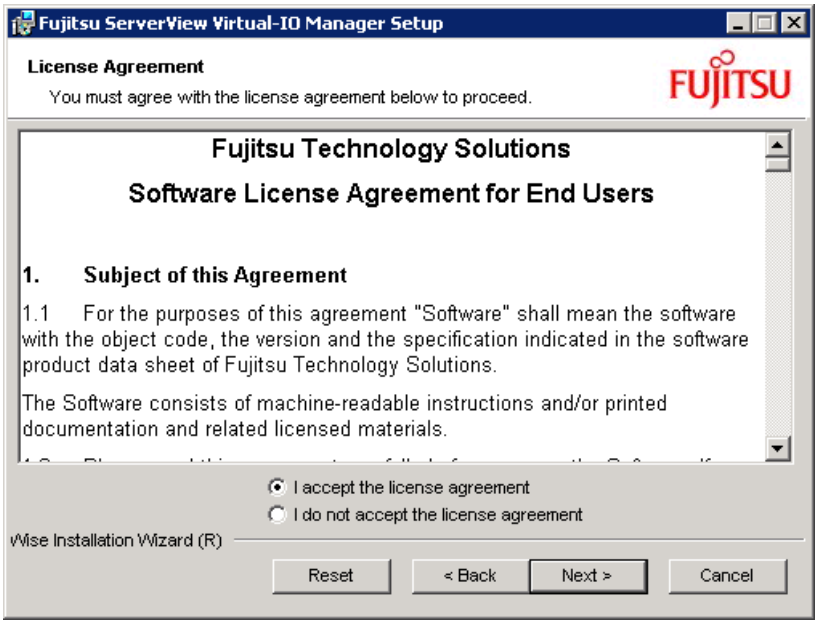

Accept the license agreement by selecting the corresponding option.

8. Click **Next**.

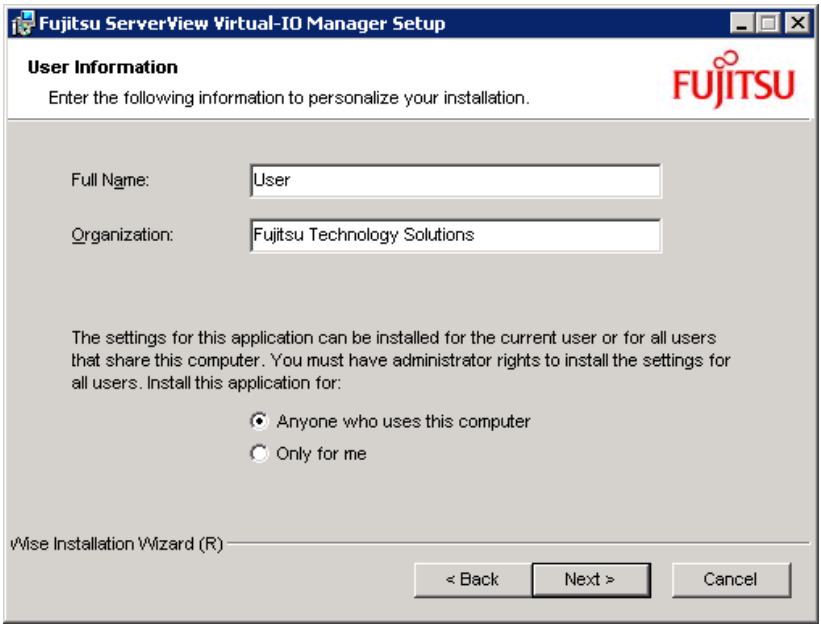

Enter your name and the name of your company/organization. You must also specify whether the settings should only apply for the current user or for any user working on this system. Select the corresponding option.

9. Click **Next**.

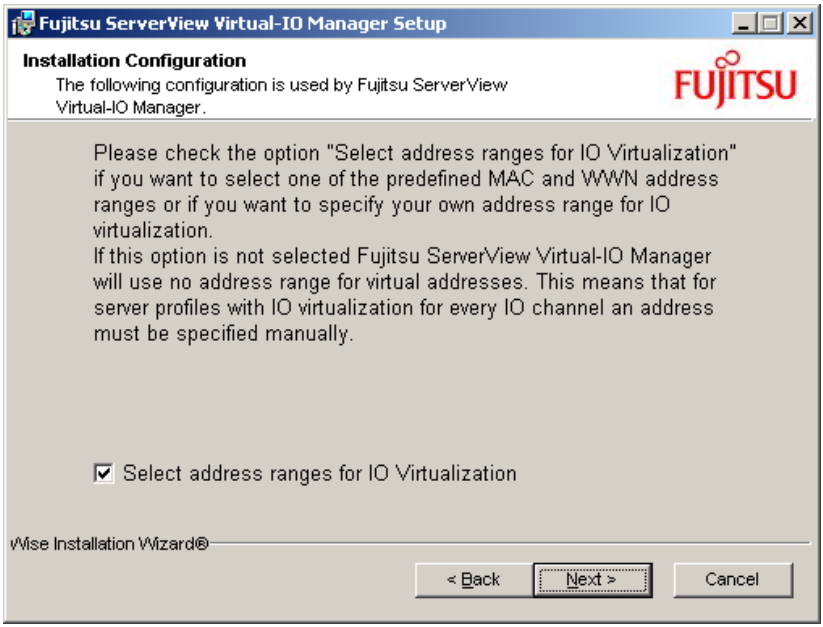

If you select **Select address ranges for IO Virtualization**, you can specify address ranges for virtual addressing.

If you want to use the automatic assignment of virtual addresses in server profiles, you must have already defined address ranges here.

10. Click **Next**.

If you did not select **Select address ranges for IO Virtualization** in the previous window, clicking **Next** brings you to the **Ready to Install the Application** window in which you start the installation.

If you selected **Select address ranges for IO Virtualization**, the following window opens:

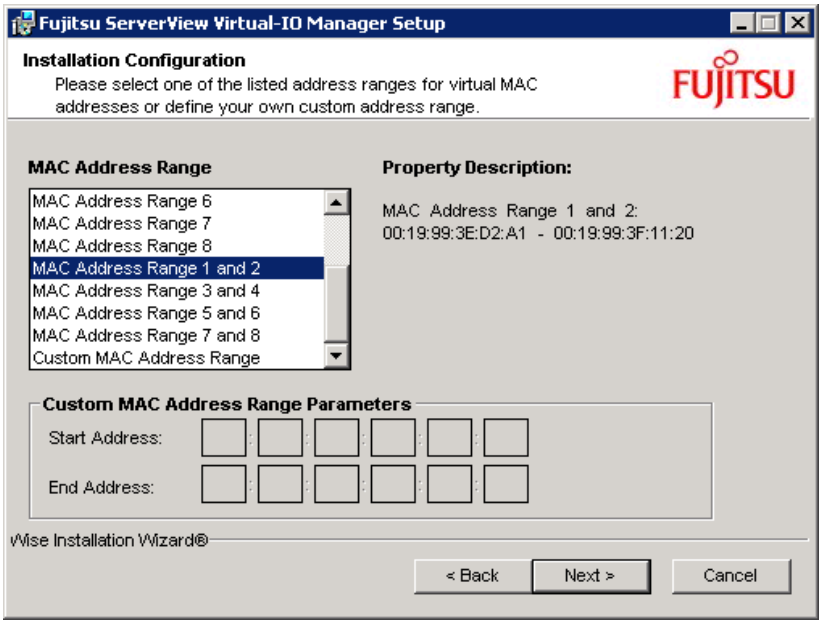

In this window, specify which address range the Virtual-IO Manager should use for virtual MAC addresses. The virtual MAC addresses are assigned automatically in a profile for the LAN ports during the server profile definition.

Eight predefined MAC address ranges are available for selection, which do not overlap (**MAC Address Range 1** to **MAC Address Range 8**). Each of these address ranges contains 8000 MAC addresses. If such a range is insufficient, you can also select a range double the size using **MAC Address Range 1 and 2** to **MAC Address Range 7 and 8**. Each of these areas contains 16000 MAC addresses.

If you have an address range of your own that you wish to use for virtual MAC addresses, then select it in the **Custom MAC Address Range** drop-down menu. In this case, the fields in **Custom MAC Address Range Parameters**, in which you can enter the start and end MAC address of the address range, become active.

The validity of the MAC address input is not checked. Please confirm that your input address is valid before you click **Next**.

If you have several installations of the Virtual-IO Manager in your LAN network, then you must ensure that the address ranges used do not overlap. Otherwise addresses may be assigned several times.

To change the range of the address after installation is finished, you must uninstall the Virtual-IO Manager and install it again.

Therefore, if it is possible that addresses may have to be added after the Virtual-IO Manager is set up, we recommend that you install it without selecting the address range. You should input a virtual address when creating the server profile.

11. Click **Next**.

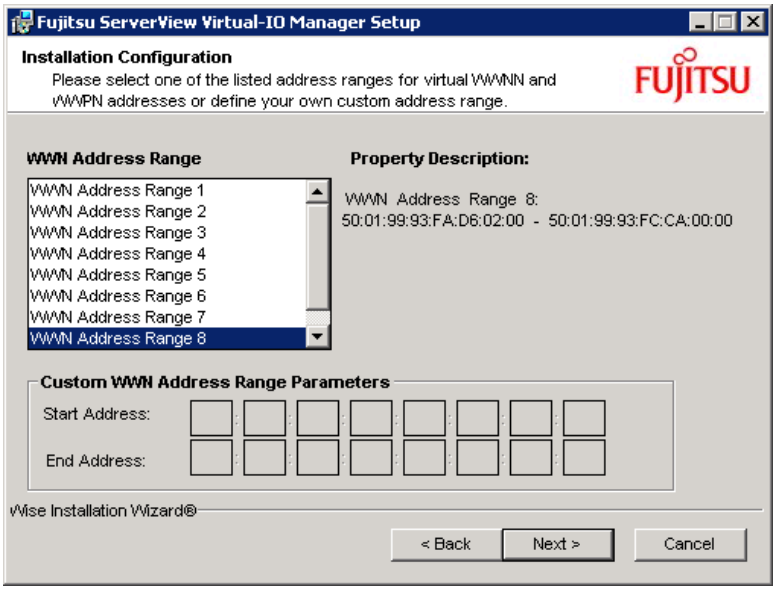

In this window, specify which address range the Virtual-IO Manager should use for virtual WWN addresses. The virtual WWN addresses are assigned automatically for the Fibre Channel ports of an optional Fibre Channel mezzanine card during the server profile definition, whereby each port has two addresses (a WWPN - World Wide Port Name and a WWNN - World Wide Node Name).

Eight predefined WWN address ranges are available for selection, which do not overlap (**WWN Address Range 1** to **WWN Address Range 8**). Each individual address range contains 32,767,487 WWN addresses.

If you have an address range of your own that you wish to use for virtual WWN addresses, then select it in the **Custom WWN Range** drop-down menu. In this case, the fields in **Custom WWN Address Range Parameters**, in which you can enter the start and end WWN address, become active.

The validity of the WWN address input is not checked. Please confirm that your input address is valid before you click **Next**.

If you have several installations of the Virtual-IO Manager in your Storage network, then you must ensure that the address ranges used do not overlap. Otherwise addresses may be assigned several times.

To change the range of the address after installation is finished, you must uninstall the Virtual-IO Manager and install it again.

Therefore, if it is possible that addresses may have to be added after the Virtual-IO Manager is set up, we recommend that you install it without selecting the address range. You should input a virtual address when creating the server profile.

12. Click **Next**.

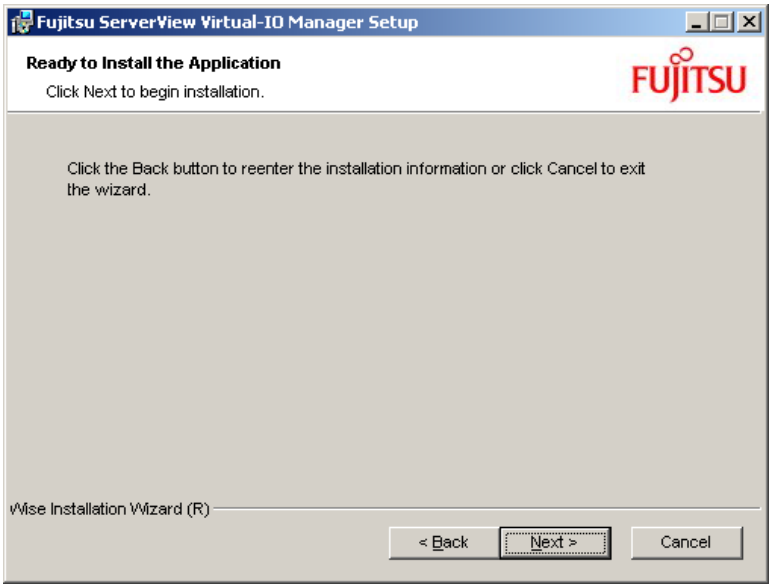

Once you have made all your entries, click **Next** to start the installation. If you want to make further changes, click **Back** to return to the previous window.

13. Click **Next**. The installation of the Virtual-IO Manager is started. The following window is then displayed:

14. Click **Next** to launch the License Manager.

15. Click **Register new license**.

Enter at least one valid license here so that you can use the Virtual-IO Manager functions. You can enter several licenses here. For more information on the License Manager, see section "License management" on page 90

- The licenses are not version bound.
- Licenses purchased with Virtual-IO Manager versions prior to V2.4 are also still valid. These licenses (v1) contain a chassis count which is multiplied by 18 to get the assign count used in licenses (v2) purchased with Virtual-IO Manager V2.4 or later.

16. Click **Exit** in the **VIOM License Manager** dialog box to exit the License Manager. The dialog box closes.
17. Click **Finish** to end the installation.

3.2.2 Installing the Virtual-IO Manager using the command line interface

The Virtual-IO Manager can be installed using the command line interface. You start the installation via the installation package **SV_VIOM.exe** which you will find on the ServerView Suite DVD 1 under **ServerView – Virtual-IO Manager**.

The installation parameters like address ranges and VIOM license can be specified by command line. The following syntax is supported:

```
SV_VIOM.exe /q [DO_ADDRESS_RANGE_SELECTION=true|false][MAC_
RANGE=<mac_range>][MAC_START="<mac_address_start>"]
[MAC_END="<mac_address_end>"][WWN_RANGE=<wwn_range>][WWN_
START="<wwn_address_start>"][WWN_END="<wwn_address_
```

```
end>"][DEBUG_MODE=true|false][VIOM_LICENSE_KEY = "<key_
value>"]
```

Command line parameters for installation:

DO_ADDRESS_RANGE_SELECTION

Possible values are:

true

You want to set MAC_RANGE and WWN_RANGE (default value).

false

You do not want to set MAC_RANGE and WWN_RANGE.

MAC_RANGE

Possible values are:

NONE

No predefined MAC address range (default value)

MAC1

Range 00:19:99:3E:D2:A1 - 00:19:99:3E:F1:E0

MAC2

Range 00:19:99:3E:F1:E1 - 00:19:99:3F:11:20

MAC3

Range 00:19:99:3F:11:21 - 00:19:99:3F:30:60

MAC4

Range 00:19:99:3F:30:61 - 00:19:99:3F:4F:A0

MAC5

Range 00:19:99:3F:4F:A1 - 00:19:99:3F:6E:E0

MAC6

Range 00:19:99:3F:6E:E1 - 00:19:99:3F:8E:20

MAC7

Range 00:19:99:3F:8E:21 - 00:19:99:3F:AD:60

MAC8
Range 00:19:99:3F:AD:61 - 00:19:99:3F:CC:A1

MAC12
Range 00:19:99:3E:D2:A1 - 00:19:99:3F:11:20

MAC34
Range 00:19:99:3F:11:21 - 00:19:99:3F:4F:A0

MAC56
Range 00:19:99:3F:4F:A1 - 00:19:99:3F:8E:20

MAC78
Range 00:19:99:3F:8E:21 - 00:19:99:3F:CC:A1

MAC_CUSTOM
Custom range must be set with the **MAC_START** and **MAC_END** parameters.

MAC_START, MAC_END
These parameters must be set if **MAC_Custom** is specified.

The values must be in hexadecimal format, for example "11:22:33:44:55:66".

WWN_RANGE
Possible values are:

NONE
No predefined WWN address range (default value)

WWN1
Range 50:01:99:93:ED:2A:10:00 -- 50:01:99:93:EF:1E:0D:FF

WWN2
Range 50:01:99:93:EF:1E:0E:00 - 50:01:99:93:F1:12:0B:FF

WWN3
Range 50:01:99:93:F1:12:0C:00 - 50:01:99:93:F3:06:09:FF

WWN4
Range 50:01:99:93:F3:06:0A:00 - 50:01:99:93:F4:FA:07:FF

WWN5
Range 50:01:99:93:F4:FA:08:00 - 50:01:99:93:F6:EE:05:FF

WWN6
Range 50:01:99:93:F6:EE:06:00 - 50:01:99:93:F8:E2:03:FF

WWN7
Range 50:01:99:93:F8:E2:04:00 - 50:01:99:93:FA:D6:02:FF

WWN8
Range 50:01:99:93:FA:D6:02:00 - 50:01:99:93:FC:CA:00:00

WWN_CUSTOM
Custom range must be set with the **WWN_START** and **WWN_END** parameters.

WWN_START, WWN_END
These parameters must be set if **WWN_Custom** is specified.

The values must be in hexadecimal format, for example "11:22:33:44:55:66:77:88".

DEBUG_MODE
Possible values are:

true
Enable debug mode (default value)

false
Disable debug mode

VIOM_LICENSE_KEY
The license key

Example

1. `SV_VIOM.exe /q MAC_RANGE=MAC_CUSTOM MAC_START="11:22:33:44:55:66" MAC_END="22:33:44:55:66:77" WWN_RANGE=WWN_CUSTOM WWN_START="33:44:55:66:77:88:99:AA" WWN_END="44:55:66:77:88:99:AA:BB" VIOM_LICENSE_KEY=abcdef`
2. `SV_VIOM.exe /q MAC_RANGE=MAC78 WWN_RANGE=WWN8 VIOM_LICENSE_KEY=abcdef`
3. `SV_VIOM.exe /q DO_ADDRESS_RANGE_SELECTION=false VIOM_LICENSE_KEY=abcdef`

3.3 Updating the Virtual-IO Manager on a Windows-based CMS

If you have already installed a previous version, an update installation runs automatically when you install the Virtual-IO Manager. In this case, all user-specific configurations and definitions remain the same.

An update installation starts in the same way as a full installation (see section "Installing the Virtual-IO Manager on a Windows-based CMS" on page 58). But the update installation takes place once you have confirmed the license agreement and exited the readme window by clicking **Next**.

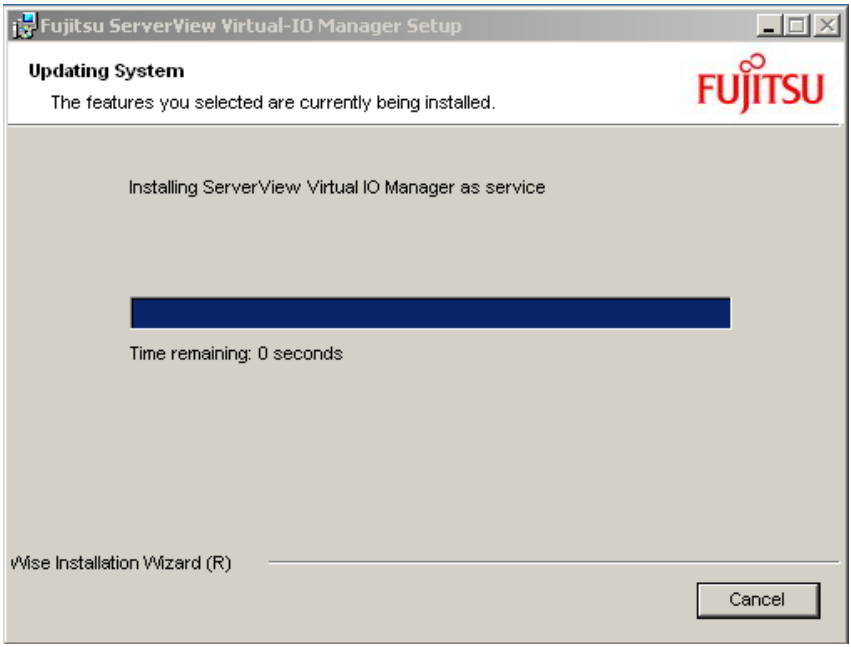

Once the update installation is complete, the final window of the installation wizard confirms that the update installation has been successful, just like in the full installation. Exit the installation wizard by clicking **Finish**.

3.4 Installing the Virtual-IO Manager on a Linux-based CMS

The corresponding software is supplied with the PRIMERGY ServerView Suite DVD1. You can find the entire software for the PRIMERGY ServerView Suite under **ServerView Software Product Selection**. To find the Virtual-IO Manager software package like **3.1.0_2012.08.15.zip** in this product selection, choose **ServerView – Virtual-IO Manager**.

Please check the requirements for installing the Virtual-IO Manager to see whether the Linux distribution you use is supported by the Virtual-IO Manager (see section "Prerequisites for the VIOM installation" on page 57).

3.4.1 Installing the Virtual-IO Manager using a graphical interface

If you want to use the graphical installation, an X Windows server should be installed on your desktop computer. Check the prerequisites for the VIOM installation (see section "Prerequisites for the VIOM installation" on page 57).

Installation process

1. Set the **DISPLAY** environment variable:

   ```
   export DISPLAY=<IP-address|host name>:0.0
   ```

   Example

   ```
   export DISPLAY=111.22.33.115:0.0
   ```

2. If you prefer to have a dedicated directory for each VIOM package, create one (e.g. **/root/VIOM_3.1.0**) and copy the zip file (e.g. **VIOM_ 3.1.0.2012.08.15.zip**) from the installation medium to that directory and unzip it:

   ```
   mkdir /root/VIOM_3.1.0
   cp VIOM_3.1.0.2012.08.15.zip /root/VIOM_3.1.0
   cd /root/VIOM_3.1.0
   unzip VIOM_3.1.0.2012.08.15.zip
   ```

3. After these preparations, start the installation by launching the GUI:

   ```
   sh install_viom.sh
   ```

   The welcome window opens.

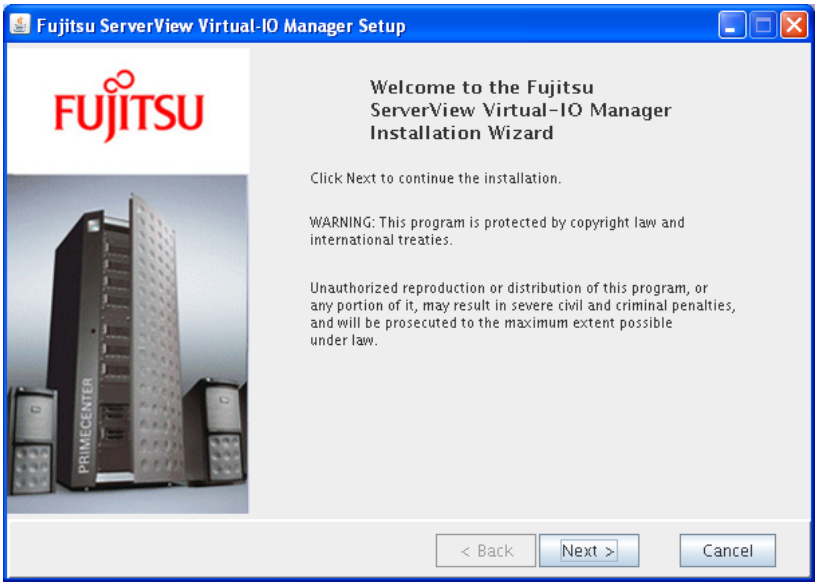

4. Click **Next**.

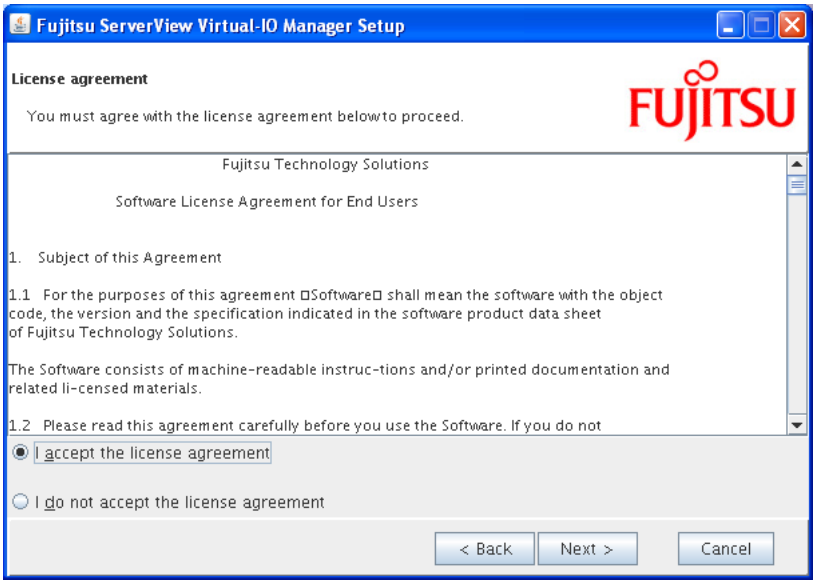

Accept the license agreement by selecting the corresponding option.

5. Click **Next**.

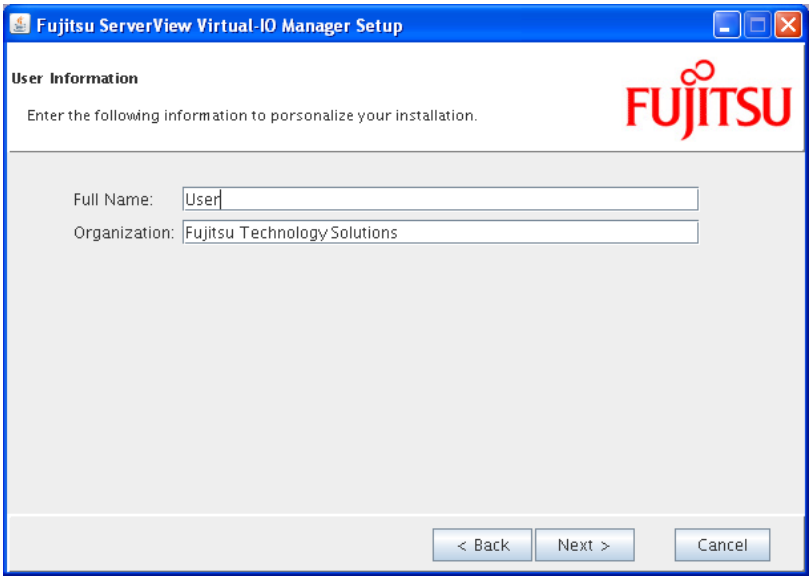

Enter your name and the name of your company/organization.

6. Click **Next**.

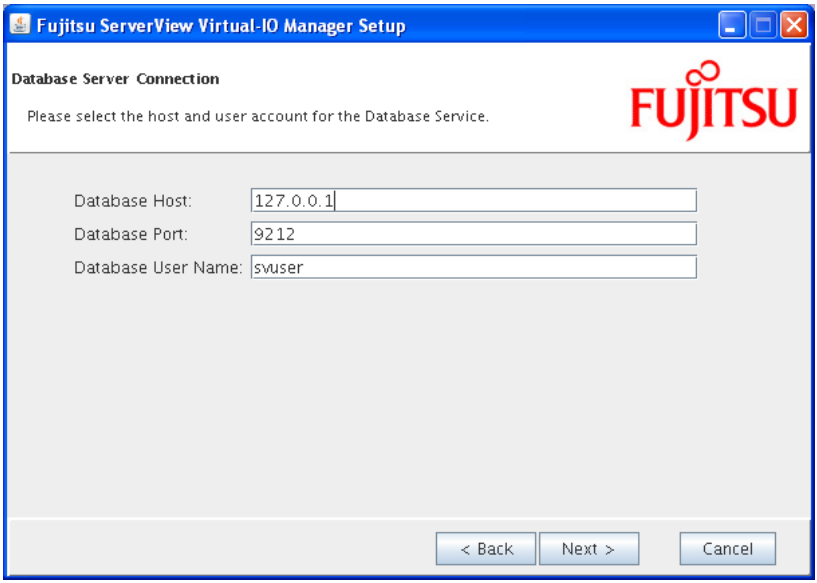

Enter the data on the ServerView database. The database user **svuser** has been created during ServerView installation and is used by VIOM as well.

7. Click **Next**.

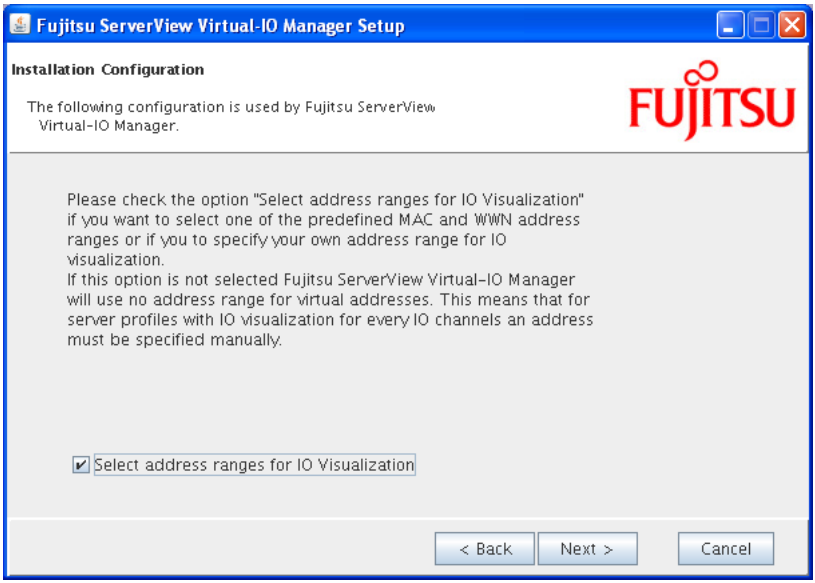

If you select **Select address ranges for IO Virtualization**, you can specify address ranges for virtual addressing.

If you want to use the automatic assignment of virtual addresses in server profiles, you must have already defined address ranges here.

8. Click **Next**.

If you did not select **Select address ranges for IO Virtualization** in the previous window, clicking **Next** brings you to the **Ready to Install the Application** window in which you start the installation.

If you selected **Select address ranges for IO Virtualization**, the following window opens:

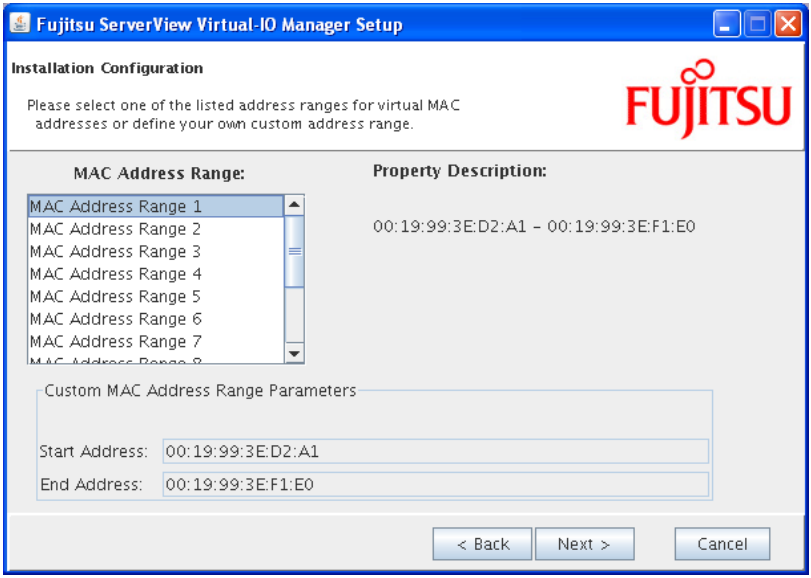

In this window, specify which address range the Virtual-IO Manager should use for virtual MAC addresses. The virtual MAC addresses are assigned automatically in a profile for the LAN ports during the server profile definition.

Eight predefined MAC address ranges are available for selection, which do not overlap (**MAC Address Range 1** to **MAC Address Range 8**). Each of these address ranges contains 8000 MAC addresses. If such a range is insufficient, you can also select a range double the size using **MAC Address Range 1 and 2** to **MAC Address Range 7 and 8**. Each of these areas contains 16000 MAC addresses.

If you have an address range of your own that you wish to use for virtual MAC addresses, then select it in the **Custom MAC Address Range** drop-down menu. In this case, the fields in **Custom MAC Address Range Parameters**, in which you can enter the start and end MAC address of the address range, become active.

Validity of MAC address input is not checked. Please advance to the next screen after confirming your input address is valid.

The validity of the MAC address input is not checked. Please confirm that your input address is valid before you click **Next**.

If you have several installations of the Virtual-IO Manager in your LAN network, then you must ensure that the address ranges used do not overlap. Otherwise addresses may be assigned several times.

To change the range of the address after installation is finished, you must uninstall the Virtual-IO Manager and install it again.

Therefore, if it is possible that addresses may have to be added after the Virtual-IO Manager is set up, we recommend that you install it without selecting the address range. You should input a virtual address when creating the server profile.

9. Click **Next**.

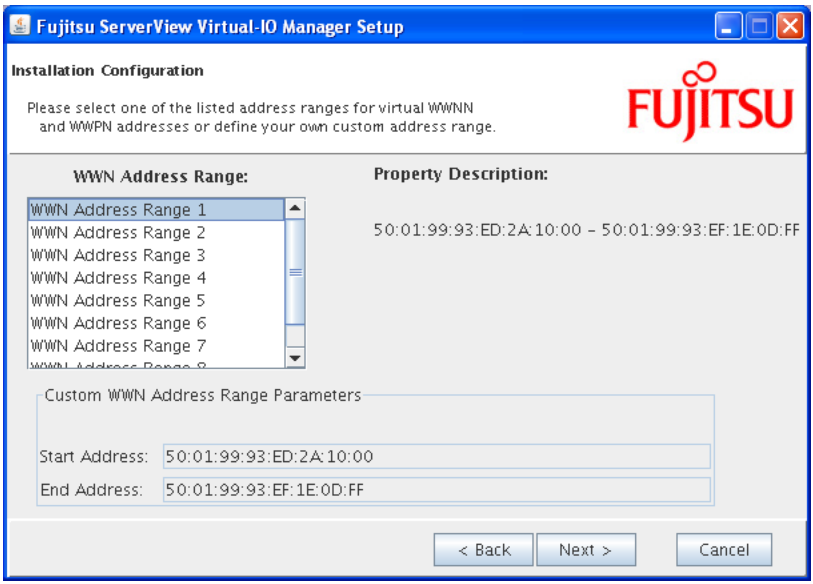

In this window, specify which address range the Virtual-IO Manager should use for virtual WWN addresses. The virtual WWN addresses are assigned automatically for the Fibre Channel ports of an optional Fibre Channel mezzanine card during the server profile definition, whereby each port has two addresses (a WWPN - World Wide Port Name and a WWNN - World Wide Node Name).

Eight predefined WWN address ranges are available for selection, which do not overlap (**WWN Address Range 1** to **WWN Address Range 8**). Each individual address range contains 32,767,487 WWN addresses.

If you have an address range of your own that you wish to use for virtual WWN addresses, then select it in the **Custom WWN Range** drop-down menu. In this case, the fields in **Custom WWN Address Range Parameters**, in which you can enter the start and end WWN address, become active.

The validity of the WWN address input is not checked. Please confirm that your input address is valid before you click **Next**.

If you have several installations of the Virtual-IO Manager in your Storage network, then you must ensure that the address ranges used do not overlap. Otherwise addresses may be assigned several times.

To change the range of the address after installation is finished, you must uninstall the Virtual-IO Manager and install it again.

Therefore, if it is possible that addresses may have to be added after the Virtual-IO Manager is set up, we recommend that you install it without selecting the address range. You should input a virtual address when creating the server profile.

10. Click **Next** to launch the License Manager.

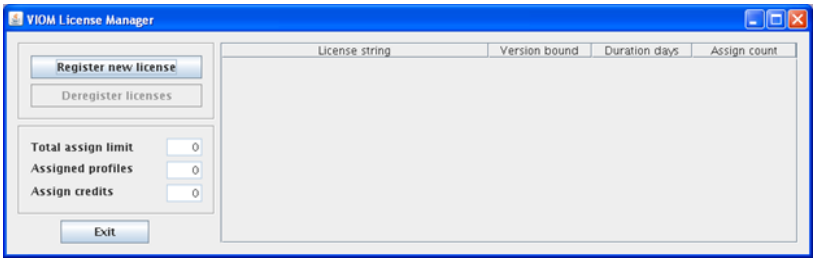

11. Click **Register new license**.

Enter at least one valid license here so that you can use the Virtual-IO Manager functions. You can enter several licenses here. For more information on the License Manager, see section "License management" on page 90. Then click **OK**

- The licenses are not version bound.
- Licenses purchased with Virtual-IO Manager versions prior to V2.4 are also still valid. These licenses (v1) contain a chassis count which is multiplied by 18 to get the assign count used in licenses (v2) purchased with Virtual-IO Manager V2.4 or later.

12. Click **Exit** in the **VIOM License Manager** dialog box to exit the License Manager. The dialog box closes.

13. In the final screen you may view the release notes or the installation log.

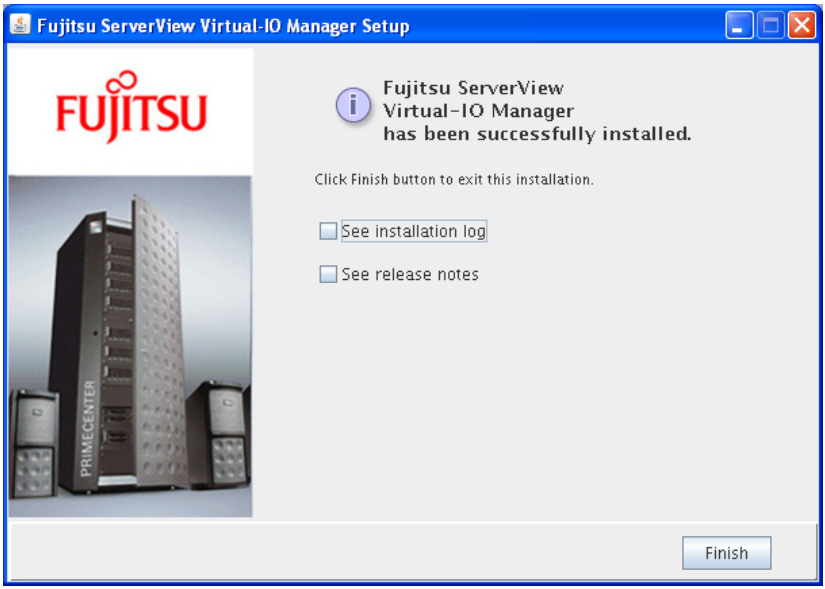

14. Click **Finish** to end the installation.

3.4.2 Installing the Virtual-IO Manager using the command line

The installation script **install_viom.sh** may also be used in a non-graphical installation, as explained here. You find the script in the directory where you unzipped the VIOM software package (see "Installing the Virtual-IO Manager using a graphical interface" on page 74). Simply use the **--no-gui** switch and enter suitable parameters.

The following syntax is supported:

```
sh install_viom.sh
[ --dbhost=<ip-address|hostname>][ --dbport=<port-number>]
[--dbuser=<username>][ --dbpasswd=<password>]
[ --mac-range=<first-mac-addr>,<last-mac-addr>|MAC<x>|MAC<xy>]
[ --wwn-range=<first-wwn-addr>,<last-wwn-addr>|WWN<x>]
[ --no-gui][ --license-key=<key-string>]
```

```
[ --check-prerequisites][ --lang[ en|ja]][ -u|--upgrade]
[ -q|--quiet][ -f|--force][ -i|--installdir][ -v|--verbose]
[ -h|--help]
```

Command line parameters for installation:

--dbhost
IP-address or hostname of the database server

--dbport
Port number of the database server

--dbuser
ServerView Operations Manager username in the database, by default **svuser**

--dbpasswd
Password in the database (by default the ServerView Operations Manager database has no password)

--mac-range
MAC address range.

Custom range must be set with the MAC start address and the MAC end address. The values must be in hexadecimal format, for example "11:22:33:44:55:66".

Predefined MAC address ranges must be separated by a comma. Possible values for predefined MAC address ranges are:

MAC1
Range 00:19:99:3E:D2:A1 - 00:19:99:3E:F1:E0

MAC2
Range 00:19:99:3E:F1:E1 - 00:19:99:3F:11:20

MAC3
Range 00:19:99:3F:11:21 - 00:19:99:3F:30:60

MAC4
Range 00:19:99:3F:30:61 - 00:19:99:3F:4F:A0

MAC5
Range 00:19:99:3F:4F:A1 - 00:19:99:3F:6E:E0

MAC6
Range 00:19:99:3F:6E:E1 - 00:19:99:3F:8E:20

MAC7
Range 00:19:99:3F:8E:21 - 00:19:99:3F:AD:60

MAC8
Range 00:19:99:3F:AD:61 - 00:19:99:3F:CC:A1

MAC12
Range 00:19:99:3E:D2:A1 - 00:19:99:3F:11:20

MAC34
Range 00:19:99:3F:11:21 - 00:19:99:3F:4F:A0

MAC56
Range 00:19:99:3F:4F:A1 - 00:19:99:3F:8E:20

MAC78
Range 00:19:99:3F:8E:21 - 00:19:99:3F:CC:A1

--wwn-range
WWN address range.

The values for a custom address range must be in hexadecimal format, for example "11:22:33:44:55:66:77:88".

Predefined WWN address ranges must be separated by a comma. Possible values for predefined WWN address ranges are:

WWN1
Range 50:01:99:93:ED:2A:10:00 -- 50:01:99:93:EF:1E:0D:FF

WWN2
Range 50:01:99:93:EF:1E:0E:00 - 50:01:99:93:F1:12:0B:FF

WWN3
Range 50:01:99:93:F1:12:0C:00 - 50:01:99:93:F3:06:09:FF

WWN4
Range 50:01:99:93:F3:06:0A:00 - 50:01:99:93:F4:FA:07:FF

WWN5
Range 50:01:99:93:F4:FA:08:00 - 50:01:99:93:F6:EE:05:FF

WWN6
Range 50:01:99:93:F6:EE:06:00 - 50:01:99:93:F8:E2:03:FF

WWN7
Range 50:01:99:93:F8:E2:04:00 - 50:01:99:93:FA:D6:02:FF

WWN8
Range 50:01:99:93:FA:D6:02:00 - 50:01:99:93:FC:CA:00:00

--no-gui
Run installation in console mode, non-gui installation

--license-key
The license key for the Virtual IO Manager

--check-prerequisites
Only check prerequisites and exit

--lang
VIOM installation language

Possible values:

en
English

ja
Japanese

--upgrade|-u
Run installation in upgrade mode if you upgrade from one VIOM version to another

--quiet|-q
Silent installation

--force|-f
Force installation without rpm and error checks

--installdir|-i
Path to the installation directory tree up to and including subdirectory **RPMS** (default:current directory)

--verbose|-v
More logging details in the log files

Example

```
sh install_viom.sh --no-gui
--dbhost=127.0.0.1 --dbport=9212
 --mac-range=00:19:99:3E:D2:A1,00:19:99:3E:F1:E0
--wwn-range=50:01:99:93:ED:2A:10:00,50:01:99:93:EF:1E:0D:FF
--license-key=AAA-BBBB-CCCC-DDDD-EEEE-FFFF-GGGG
```

3.4.3 Important directories of Virtual-IO Manager

/opt/fujitsu/ServerViewSuite/plugins/viom
In all Linux distributions, this is the installation directory where the binaries and libraries are stored.

/var/fujitsu/ServerViewSuite/viom
Directory in which variable information of VIOM is usually stored.

/var/log/fujitsu/ServerViewSuite/viom
Directory in which the VIOM log files are stored.

3.4.4 Collecting diagnostic information

Sometimes it might be necessary to gather diagnostics information in order to send it to your Fujitsu support service.

In the **Manager** subdirectory of the VIOM installation directory, you will find a **dump.sh** script that you may launch using the following command:

```
cd /opt/fujitsu/ServerViewSuite/plugins/viom/Manager
sh dump.sh
```

The **dump.sh** script collects and generates a set of diagnostics files, zips them and stores the result in the directory:

/var/fujitsu/ServerViewSuite/viom/dumps/

3.5 Updating the Virtual-IO Manager on a Linux-based CMS

If you have already installed a previous version, you can run an update installation. In this case, you must specify the **–upgrade** option when invoking the installation script:

```
sh install_viom.sh -upgrade
```

When running an update installation, all user-specific configurations and definitions remain the same.

An update installation starts in the same way as a full installation (see section "Installing the Virtual-IO Manager on a Linux-based CMS" on page 73). But the update installation takes place once you have confirmed the license agreement and exited the readme window by clicking **Next**.

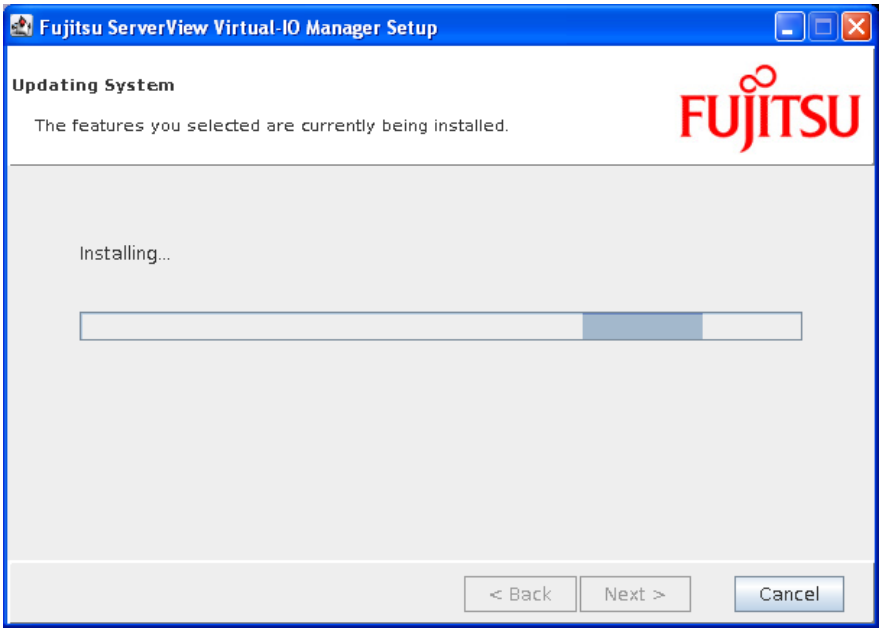

Once the update installation is complete, the final window of the installation wizard confirms that the update installation has been successful, just like in the full installation. Exit the installation wizard by clicking **Finish**.

3.6 License management

You need at least one license in order to use VIOM. You purchase licenses for a certain number of server profiles that can be assigned to server blades using VIOM.

- The licenses are not version bound.
- Registered licenses for an installed Virtual-IO Manager version prior to V2.4 are multiplied internally by 18. This means that the chassis count of these licenses (v1) is converted (x18) into the assign count used in licenses (v2) used as of Virtual-IO Manager V2.4.

The specified licenses are verified when you launch VIOM:

- If you do not have a valid license, you cannot log in to the VIOM Manager.
- If the number of currently assigned server profiles exceeds the number of server profiles you are allowed to assign, you cannot make any further configuration changes except for **Unassign Profile** (see section "Deleting profile assignments" on page 298) or **Unmanage**(see section "Deactivating management with VIOM" on page 261).

A license check is performed automatically when you install VIOM (see section "Installing the Virtual-IO Manager on a Windows-based CMS" on page 58) and periodically during the running time of the VIOM service.

On Windows, you can also launch the License Manager, the tool for managing licenses, separately from the Windows Start Menu, e.g. to specify new or additional licenses.

On Linux, you can also launch the License Manager by starting the **license.sh** script which you find in the installation directory of VIOM.

1. On Windows: Choose **Start – [All] Programs – Fujitsu – ServerView Suite – Virtual-IO Manager – License Management**.

   On Linux: Execute the command **sh license.sh**.

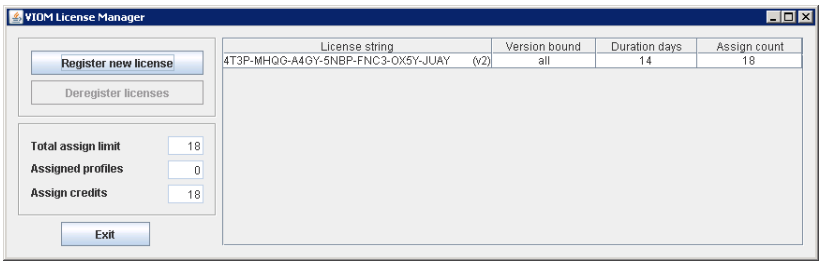

   Invalid licenses (e.g. expired or banned licenses) are shown with a red background.

2. Click **Register new license**.

   Enter a valid license here. You can also enter several licenses by clicking **Register new license** again. Each license contains a number of assigns. This license permits you to assign this number of server profiles.

   Licenses purchased with Virtual-IO Manager versions prior to V2.4 are also still valid. These licenses (v1) contain a chassis count which is multiplied by 18 to get the assign count used in licenses (v2) purchased with Virtual-IO Manager V2.4 or later.

3. Click **Exit** in the **VIOM License Manager** dialog box to exit the License Manager.

On Linux you can also add a licence without graphical user interface:

```
cd /opt/fujitsu/ServerViewSuite/plugins/viom/Manager
sh license.sh --license-key=<key> --no-gui
```

You can also delete licenses, provided that the total number of assignments of the remaining licenses is not smaller than the number of server profiles assigned. To do this, launch the License Manager from the Windows Start Menu or with the **license.sh** script.

1. On Windows: Choose **Start – [All] Programs – Fujitsu – ServerView Suite – Virtual-IO Manager – License Management**.

   On Linux: Execute the command **sh license.sh**.

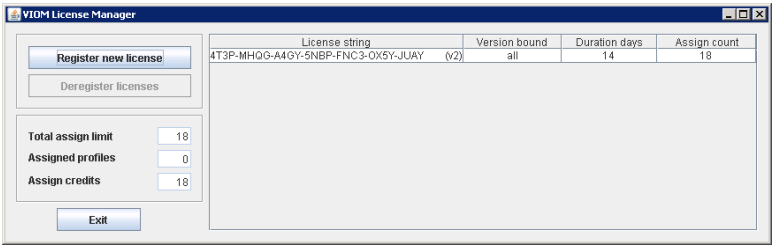

2. Select the relevant license and click **Deregister licenses**. The selected license is deleted. If the number of allowed assignments is smaller than the number of assigned server profiles, an error message is displayed and the deletion process canceled.

On Linux you can also add a licence without graphical user interface:

```
cd /opt/fujitsu/ServerViewSuite/plugins/viom/Manager
java -jar ./VIOM-LICENSE-MANAGER.jar -remove <key>
```

In addition to using the License Manager to manage licenses, you can display the licenses assigned in the Virtual-IO Manager interface. To do this, you must click the **Show Licenses** button on the **Virtual-IO Manager** tab.

You use the **Show Licenses** button on the **Virtual-IO Manager** tab to request information on the licenses assigned.

1. To do this, click the **Show Licenses** button on the **Virtual-IO Manager** tab.
2. The **Licenses Information** dialog box opens.

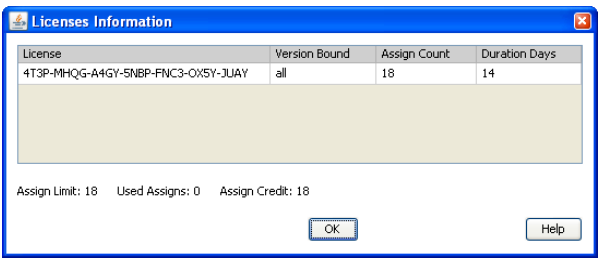

3. The information displayed includes the following:
   - License code
   - Validity period (only relevant for demo licenses)
   - Max. number of assignments that are allowed with the licenses
   - Number of currently assigned server profiles
   - Number of server profiles that can still be assigned with the licenses

For more information, see section "Displaying license information" on page 275.

3.7 Updating ServerView Operations Manager

ServerView Operations Manager may be updated at any time while the Virtual-IO Manager is idle. Thus you should stop the Virtual-IO Manager Service before starting the update.

Updating ServerView Operations Manager consists of three steps:

- Perform an update installation of ServerView Operations Manager
- Check that the updated version works okay
- Start the Virtual-IO Manager Service

After performing a modify installation of ServerView Operations Manager, you need to start the Virtual-IO Manager Service.

When you need to uninstall existing ServerView Operations Manager and install the new one, you should uninstall existing Virtual-IO Manager:

- Save your configuration by using the **Configuration Backup / Restore** button on the **Virtual-IO Manager** tab
- Uninstall Virtual-IO Manager
- Uninstall ServerView Operations Manager
- Install ServerView Operations Manager
- Install Virtual-IO Manager
- Restore your configuration by using the **Configuration Backup / Restore** button on the **Virtual-IO Manager** tab

3.8 Upgrading or moving the SQL Server database

The ServerView Operations Manager manual (Installation under Windows) explains what you must be aware of when updating your SQL Server database on the central management station (CMS) or on a remote node.

The following is based on that information. In chapter "Upgrade installations of SQL Server" two major cases are distinguished:

- Upgrading an SQL Server instance while keeping the instance name
- Moving an SQL Server instance to an instance with a different name

In the former case there there is one thing to add if the Virtual-IO Manager is running on the central management station. Just make sure that both the Virtual-IO Manager and the Backup Service are stopped before you start upgrading the SQL Server.

In the latter case you must perform two additional steps (while the VIOM services are stopped)

- When you have created the ServerViewDB database in the new instance, execute a similar command sequence for ViomDB

```
CREATE DATABASE ViomDB
ON PRIMARY (FILENAME='<data_path>\ViomDB.mdf')
```

```
LOG ON (FILENAME='<data_path>\ViomDB_log.LDF')
FOR ATTACH
DBCC UPDATEUSAGE('ViomDB')
```

- In the installation directory of the Virtual-IO Manager you will find a Java properties file called **ViomConfig.properties**. Change the instance name to the new value, e.g.:

```
SqlServerInstance=(local)\\SQLSERVERVIEW2
```

3.9 Uninstalling the Virtual-IO Manager

Before you uninstall VIOM, you should unmanage all servers. After uninstalling, all information will be lost. Therefore save your server profiles if you want to use them later.

3.9.1 Uninstalling the Virtual-IO Manager on a Windows-based CMS

You uninstall the Virtual-IO Manager via the Windows start menu:

- On Windows Server 2003

  Select **Start – Settings – Control Panel – Add/Remove Programs**.

- On Windows Server 2008

  Select **Start – Control Panel – Programs and Features**. To display the product versions, select **Choose Details** from the **View** menu. Select the **Version** option and click **OK**

Under **Fujitsu ServerView Virtual-IO Manager** select the entry **Remove**. This uninstalls the complete package.

3.9.2 Uninstalling the Virtual-IO Manager on a Linux-based CMS

In order to uninstall VIOM use the **unstall_viom.sh** script which is found in the installation directory of VIOM:

```
cd /opt/fujitsu/ServerViewSuite/plugins/viom/Manager
sh uninstall_viom.sh
```

4 Configuration

In the following sections, you find information on the required configurations on the managed server.

4.1 Configurations on the managed BX600 Blade Server

In section "Supported hardware configurations for the connection blades" on page 97, you can find out what combinations of I/O connection blades are supported by VIOM.

To manage a blade server with VIOM, you must carry out the following preparatory steps:

- Configure the management blade or management blades, see section "Configuring the BX600 management blade" on page 99
- Configure the I/O connection blades, see section "Configuring the I/O connection blades" on page 101
- Define the connection of the IBP modules with the network, see section "Connecting IBP modules" on page 104
- Add the blade server to the server list of the ServerView Operations Manager, see section "Adding a server to the ServerView server list" on page 135

4.1.1 Supported hardware configurations for the connection blades

To define dedicated network connections for server blades with networks, IBP modules are required.

- At least one IBP module must be installed in fabric 1 or in fabric 2 of the blade server chassis.
- To support a Fibre Channel mezzanine card, you must install a 4 GB Brocade switch blade SW4016-D4 (Fibre Channel switch blade for BX600)

in fabric 2. FC switch blades do not support any dedicated network connections, but just the virtualization of the I/O parameters instead.

- If you only intend to work with virtual I/O parameters (Open Fabric mode) and not with dedicated network connections, you can also install non-VIOM-capable LAN modules in fabric 1 or fabric 2. Furthermore, you can operate non-VIOM-capable LAN modules in fabric 1 with FC modules in fabric 2.

4.1.1.1 LAN hardware configuration

Using IBP modules

Fabric 1 and fabric 2 have identical configuration restrictions in relation to the IBP modules. The following table shows you which configurations are supported in Fabric 1 and 2:

Bay 1 (Bay 3)	Bay 2 (Bay 4)
IBP 10/6	IBP 10/6 (optional)
IBP 30/12	IBP 30/12 (optional[1])

[1] Even though this constellation is possible, it does not make sense because only the first port can be used from the IBP 30/12 in bay 1 or bay 3.

You can also use different IBP modules within one fabric. In this constellation, however, not all downlinks can be used by the larger model.

Operation in Open Fabric mode

Open Fabric mode refers to the use of non-VIOM-capable LAN modules together with the Virtual-IO Manager. In this case, you can only work with virtual I/O parameters, and not with dedicated network connections.

Therefore, if you do not want to use any routing, all combinations are possible even with non-VIOM-capable I/O connection blades. You must take the following into account:

- You cannot mix LAN and IBP modules within a fabric.
- You can only assign profiles if no networks are defined for their ports.

  You can force the assignment of a profile with network definitions by answering the subsequent query accordingly. In doing so, however, the network definitions are ignored.

- You can operate LAN models of different sizes within a fabric, although you cannot use all download links from the larger model.

4.1.1.2 Fibre Channel hardware configuration

You can only install the Fibre Channel switch blades in fabric 2. They can be operated in two different modes:

- in normal switch mode
- in Access Gateway mode

As a result, different combination options are possible. The following table shows you which configurations are possible and which are supported in fabric 2:

Bay 3	Bay 4
Access Gateway mode	Access Gateway mode
Standard switch mode	Standard switch mode
Access Gateway mode	-------
-------	Access Gateway mode
Standard switch mode	-------
-------	Standard switch mode

4.1.2 Configuring the BX600 management blade

The blade server must be fitted with one or two S3 management blades.

The management blade has two user-friendly user interfaces: a Web interface and a Remote Manager interface per Telnet or SSH protocol (a connection to the management blade is established per Telnet or SSH). You can

find a detailed description on the interfaces in the ServerView Management Blade manual for BX600.

Check the following settings in the management blade:

- The management blade must be installed with a specific firmware version. To find out which firmware version you need, see the release notes included.

  You can check the firmware version in the Remote Manager of the management blade. To do this, choose **(1) Management Agent – (2) Management Blade**. The **Management Blade Firmware Version** parameter displays the installed firmware version. You can also check the version via the Web interface.

  If your firmware version is lower than the version required, you must update it to the required firmware version before you activate VIOM management for the blade server. You can run an update via the Remote Manager of the management blade. For a MMB S3, choose **(4) TFTP update – (6) Management Blade Update Enable**.

  After selecting **(5) TFTP Update in (4) Management Blade Image File Name**, you first specify the name of the firmware file with the new firmware version.
- You must configure the management blade and connect it to the local network so that it can be accessed from the central management station on which the ServerView Operations Manager and VIOM are installed.
- Access to the management blade via Telnet or SSH is essential for VIOM. You must configure the management blade via its Web interface. To do this, choose **Chassis Settings – Network Interface – Telnet** or **Chassis Settings – Network Interface – SSH**.
- Ensure that the **Automatic Inventory Retrieval** parameter is set to **automatic**. Then the management blade automatically generates the inventory information required by VIOM when the server blade is inserted.

  

  VIOM needs this inventory information specifically for the server blades that support the I/O virtualization.

To check the parameter setting in the Remote Manager of the management blade, choose **(1) Management Agent – (3) System Information – (11) Automatic Inventory Retrieval**.

- To ensure that the blade server names in the ServerView server list are unique, you must assign a system name to the blade server chassis. System names are assigned when the management blade is configured.

4.1.3 Configuring the I/O connection blades

The IBP modules in fabric 1 and the optional connection blades (IBP or Brocade Fibre Channel switch blades) in fabric 2 must be configured and connected to the local network so that they can be accessed from the central management station on which the ServerView Operations Manager and VIOM are installed.

You can configure the network parameters of an I/O connection blade using the easy-to-use Web interface or you can use the Remote Manager of the management blade.

Configuration using the Web interface

You can configure the network parameters of an I/O connection blade using the Web interface of the management blade.

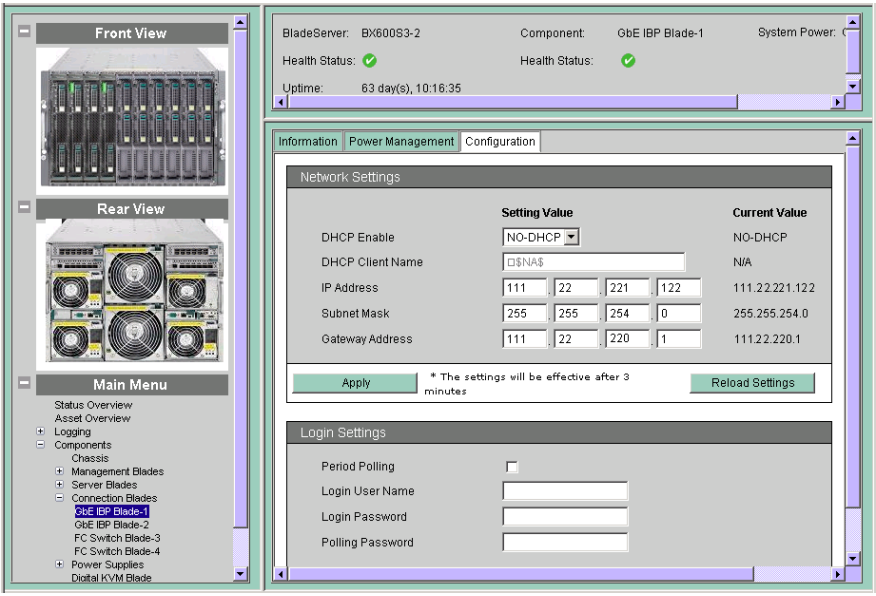

You configure the network parameters of an I/O connection blade as follows:

- In **DHCP Enable**, select whether the connection blade should be given its IP address from an DHCP server or not. If you select **NO DHCP**, you must enter the IP address, subnet mask and the gateway address.
- Click **Apply** to activate the settings.

Configuration using the Remote Manager

You can also configure the network parameters of the I/O connection blades using the Remote Manager of the management blade.

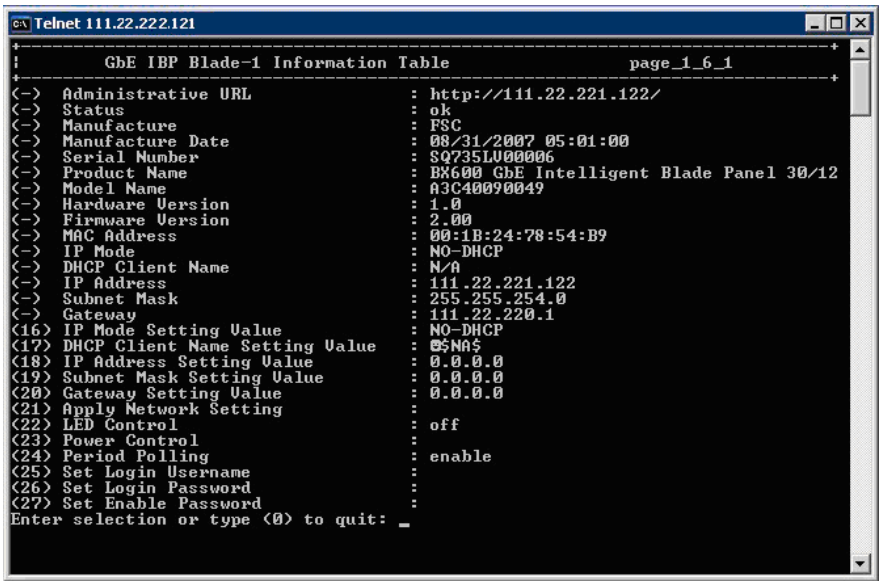

You configure the network parameters of an I/O connection blade as follows:

- In the Remote Manager, choose **(1) Management Agent – (6) Connection Blade**.
- Select an I/O connection blade.
- Make the relevant entries for the following parameters
  - **IP Mode Setting Value**
  - **IP Address Setting Value**
  - **Subnet Mask Setting Value**
  - **Gateway Setting Value**
- Activate the settings by clicking **Apply Network Setting**.

For Fibre Channel switch modules (select **FC Switch Blade**), you must also specify access data for the parameters **Set Login Username** and **Set Login Password** beforehand for the corresponding I/O connection blade in the management blade.

4.1.4 Connecting IBP modules

The IBP modules must be connected to the management LAN via the first uplink port, in other words, the first external LAN port. Here, the management LAN is the network to which the central management station is connected and via which you can access the management blades and all I/O connection blades.

The first uplink port is always the external uplink port of an IBP module on the very right of the upper row (as seen from behind).

Figure 14: First uplink port of an IBP module

For an IBP 10/6, the first uplink port is port 0/11, and for an IBP 30/12, it is port 0/31. IBP administration must be granted access to this port.

For the external switch to which an IBP module is connected for administration, you have to ensure that the Spanning Tree Protocol (STP) is switched off for the port, which is connected to the first uplink port of the IBP module.

Before you activate management of a blade server with VIOM, only this port is allowed to have a connection to external LAN switches. Other LAN cables for the connection with other networks to external switches may only be connected once the management of the blade server with VIOM has been activated. This is because the first six or first eight uplink ports form a static link aggregation group (LAG) in the IBP standard configuration.

4.1.4.1 Network - Overview

All components such as the management blades and I/O connection blades, the central management station, and the management console must be interconnected via LAN. It is recommended that you set up a management LAN that is separate from the productive LAN, as illustrated in the following figure.

If there are not enough physical ports available, a VLAN-based management LAN can also be implemented.

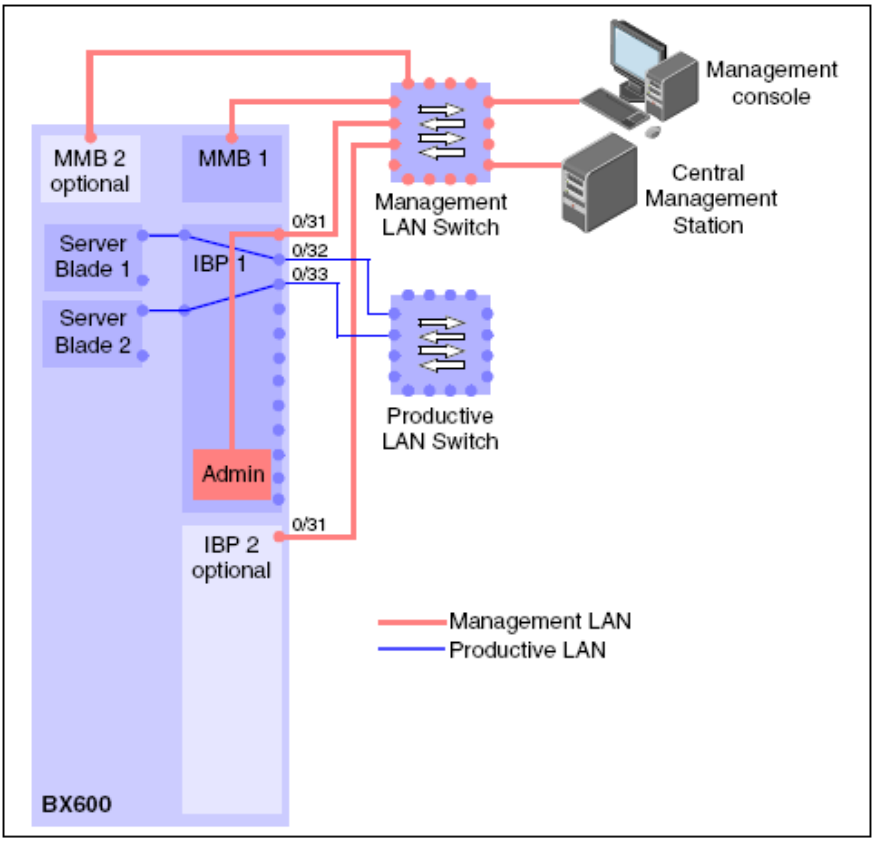

Figure 15: BX600: Network - Overview

Port 0/11 for an IBP 10/6 or port 0/31 for an IBP 30/12 must be connected to the central management station.

4.1.4.2 Notes and recommendations

We recommend that you do not use the first uplink port in the network definitions. If you cannot avoid doing so, then you must note the following:

- The speed definition used must remain set to **auto negotiate**.
- The network that uses this uplink port must not be deleted.

Otherwise, you may have a situation whereby the corresponding IBP module can no longer be accessed from the central management station. If this happens, you must use the console redirection function of the management blade to configure the IBP module so that it can be accessed again via the management LAN of the management station. You do this as follows:

- Establish a Telnet or SSH connection to the management blade of the affected blade server chassis.
  Example: **Telnet 10.10.10.50 3172**
- Enter the authentication data for the management blade.
- In the Remote Manager, choose (**3) Console Redirection – (2) Console Redirect Connection Blade**.
- Select the affected I/O connection blade (e. g. **(1) Console Redirect GbE IBP Blade 1**).
- You are now connected to the selected I/O connection blade. To access the connection blade from the central management station, the first uplink port must belong to an uplink set to which a port group is also assigned.

You can find out what the current settings are by entering the commands **show uplink-set** and **show port-group**.

Example

```
(Vty-0) #show uplink-set

Uplink   External      External      External     Link  Port   IGMP
Set Name ports         active ports  backup ports state Backup snoop LACP

-------- ------------- ------------ ------------ ----- ------ ----- ----

default  -              -             -           yes   no     yes   no

1_NET    0/11           0/11          -           yes   yes    yes   no

1_NW1    0/12,0/13,0/14 0/13,0/14     0/12        yes   yes    yes   no

1_NW2    0/15,0/16      0/15          0/16        yes   yes    yes   no

(Vty-0

(Vty-0) #show port-group

Port Group Name      Internal Ports  Uplink Set Name      External Ports

-------------------  --------------  -------------------  --------------

default              -               default              -

1_INTERN_1           -                                    -

1_INTERN_2           -                                    -
1_NET                0/1,0/2         1_NET                0/11
1_NW1                -               1_NW1                0/12,0/13,0/14
1_NW2                -               1_NW2                0/15,0/16
```

In the above example, the first uplink port 0/11 belongs to the uplink set **1_NET**, and the port group **1_NET** is assigned to this uplink set.

If port 0/11 does not belong to any uplink set, but should belong to uplink set **1_NET**, then enter the following commands:

```
(Vty-0) #configure

(Vty-0) (Config)#interface 0/11

(Vty-0) (Interface 0/11)#uplink-set 1_NET

(Vty-0) (Interface 0/11)# exit

(Vty-0) (Config)#exit

(Vty-0) #
```

If the first uplink port 0/11 belongs to the uplink set **1_NET**, but no port group

is assigned to this uplink set (because, for example, **1_NET** is to be deleted), then you can restore the LAN connection to the IBP by adding the port to the standard uplink set. Enter the following commands:

```
(Vty-0) #configure
(Vty-0) (Config)#interface 0/11
(Vty-0) (Interface 0/11)#uplink-set default
(Vty-0) (Interface 0/11)# exit
(Vty-0) (Config)#exit
(Vty-0) #
```

In these cases, however, the configuration used from VIOM no longer matches the actual configuration. In this case, you should use the VIOM user interface for the affected IBP module to restore the configuration.

- To do this, go to the **Setup** tab and click the affected IBP module in the rear view of the blade server chassis.

  The details for the IBP module and the **Restore IBP** button are then displayed at the bottom on the right.

- Click the **Restore IBP** button. The IBP module is then programmed according to the configuration saved in VIOM.

4.2 Configurations on the managed BX400 Blade Server

In section "Supported hardware configurations for the connection blades" on page 110 you can find out what combinations of I/O connection blades are supported by VIOM.

To manage a blade server with VIOM, you must carry out the following preparatory steps:

- Configure the management blade or management blades, see section "Configuring the BX400 management blade" on page 113
- Configure the I/O connection blades, see section "Configuring the I/O connection blades" on page 114

- Define the connection of the IBP modules with the network, see section "Connecting IBP modules" on page 115
- Add the blade server to the server list of the ServerView Operations Manager, see section "Adding a server to the ServerView server list" on page 135

4.2.1 Supported hardware configurations for the connection blades

The following connection blades are supported in a BX400:

- The connection blades PY CB Eth Switch/IBP 1Gb 36/8+2, PY CB Eth Switch/IBP 1Gb 36/12, PY CB Eth Switch/IBP 1Gb 18/6, and PY CB Eth Switch/IBP 10 Gb 18/8 as LAN connection blades. LAN connection blades can be inserted in connection bay 1 of fabric 1 in order to connect the onboard LAN ports, in connection bay 2 of fabric 2 in order to connect the onboard LAN ports of the first mezzanine card, or in connection bays 3 and 4 of fabric 3 in order to connect the LAN ports of the second mezzanine card.

  To define dedicated network connections for server blades with networks, you require LAN connection blades in IBP mode. The PY CB Eth Switch/IBP 1Gb 36/8+2, PY CB Eth Switch/IBP 1Gb 36/12, PY CB Eth Switch/IBP 1Gb 18/6, and PY CB Eth Switch/IBP 10 Gb 18/8 connection blades support a "normal" Layer 2 switch mode and IBP mode.

  If you only intend to work with virtual I/O parameters (Open Fabric mode of the Virtual-IO Manager) and not with dedicated network connections, you can also use the connection blades in switch mode or you can use LAN connection blades that do not support IBP mode. This applies separately to the connection blades in each fabric. In other words, you can operate one connection blade in IBP mode in fabric 1 and one connection blade in switch mode in fabric 2.
- You must install an 8 GB Brocade switch blade in fabric 2 and/or 3 in order to support a Fibre Channel mezzanine card.

4.2.1.1 LAN hardware configuration

Using IBP Modules

The following LAN connection blades in IBP mode can be used in fabric 1 (connection bay 1) and fabric 2 (connection bay 2) for the first mezzanine card:

Bay 1 (Bay2)
PY CB Eth Switch/IBP 1Gb 36/8+2 in IBP mode
PY CB Eth Switch/IBP 1Gb 36/12 in IBP mode
PY CB Eth Switch/IBP 1Gb 18/6 in IBP mode
PY CB Eth Switch/IBP 10 Gb 18/8 in IBP mode

Fabric 3 has configuration restrictions in relation to the IBP modules. The following table shows you which configurations are supported in fabric 3.

Bay 3	Bay 4
PY CB Eth Switch/IBP 1Gb 36/8+2 in IBP mode	PY CB Eth Switch/IBP 1Gb 36/8+2 in IBP mode (optional[1])
PY CB Eth Switch/IBP 1Gb 36/12 in IBP mode	PY CB Eth Switch/IBP 1Gb 36/12 in IBP mode (optional[1])
PY CB Eth Switch/IBP 1Gb 18/6 in IBP mode	PY CB Eth Switch/IBP 1Gb 18/6 in IBP mode (optional[1])
PY CB Eth Switch/IBP 10 Gb 18/8 in IBP mode	PY CB Eth Switch/IBP 10 Gb 18/8 in IBP mode (optional[1])

[1] Even though this constellation is possible, it does not make sense because only the first port can be used from the IBP module in bay 3.

Operation in Open Fabric mode

Open Fabric mode refers to the use of non-VIOM-capable LAN modules together with the Virtual-IO Manager. In this case, you can only work with virtual I/O parameters, and not with dedicated network connections.

Therefore, if you do not want to use any routing, all combinations are possible even with non-VIOM-capable I/O connection blades. You must take the following into account:

- You cannot mix LAN and IBP modules within fabric 3.
- You can only assign profiles if no networks are defined for their ports.

  You can force the assignment of a profile with network definitions by selecting the **Ignore ext. LAN connections** option in the profile selection dialog box or by answering the subsequent query accordingly. In doing so, however, the network definitions are ignored.

- You can operate LAN models of different sizes within fabric 3, although you cannot use all download links from the larger model.

4.2.1.2 Fibre Channel hardware configuration

You can install the Fibre Channel switch blades in fabric 2 or fabric 3. They can be operated in two different modes:

- in normal switch mode
- in Access Gateway mode

As a result, different combination options are possible. Fabric 2 and fabric 3 have identical configuration restrictions in relation to the Fibre Channel switch blades. The following table shows you which configurations are supported in fabric 3:

Bay 3	Bay 4
Access Gateway mode	Access Gateway mode (optional[1])
Standard switch mode	Standard switch mode (optional[1])

[1]Even though this constellation is possible, it does not make sense because only the first port can be used from the Fibre Channel switch blade in bay 3.

4.2.2 Configuring the BX400 management blade

The blade server must be fitted with at least one S1 management blade.

The BX400 management blade has a user friendly Web interface. You can find a detailed description on the interfaces in the ServerView Management Blade manual for BX400.

Check the following settings in the management blade:

- The management blade must be installed with a specific firmware version. To find out which firmware version you need, see the release notes included. You can check the firmware version in the Web interface by selecting a management blade component on the left and then checking the value of **Firmware Version** on the **Information** tab.

  If your firmware version is lower than the version required, you must update it to the required firmware version before you activate VIOM management for the blade server. You can run an update via the Web interface of the management blade. Open the **Information / Operation** section on the left side and select **Firmware Update** within **Operation**.

- You must configure the management blade and connect it to the local network so that the management LAN port of the management blade can be accessed from the central management station on which the ServerView Operations Manager and VIOM is installed.

- Access to the management blade via Telnet or SSH is essential for VIOM. You must configure the management blade via its Web interface. To do this, choose **Settings – System Unit – Network Interface – Management LAN**. On the **Ethernet** tab, check the telnet and SSH settings in the **Text based Access** section.

- Ensure that the **Automatic Inventory Retrieval** parameter is set to **automatic**, so that the management blade automatically generates the inventory information required by VIOM when the server blade is inserted.

  

  VIOM needs this inventory information specifically for the server blades that support the I/O virtualization.

To check the parameter setting in the Web interface open the **Components** section and select **System – System Unit**. The **System Information** section on the right contains the option **Automatic Inventory Retrieval**. If the value of it is not set to **automatic**, please set it to **automatic** and remove and insert all blades or perform a power-off and power-on of the blade server chassis.

- To ensure that the blade server names in the ServerView server list are unique, you must assign a system name to the blade server chassis. System names are assigned when the management blade is configured.

4.2.3 Configuring the I/O connection blades

The connection blades in fabric 1 and the optional connection blades in fabric 2 and 3 must be configured and connected to the local network so that they can be accessed from the central management station on which the ServerView Operations Manager and VIOM are installed.

You can configure the network parameters of an I/O connection blade using the easy-to-use Web interface of the management blade.

Configuration using the Web interface

It is important that the I/O connection blades of the BX400 can be accessed from the management station via their management port.

You can configure the network parameters of the management port of an I/O connection blade using the Web interface of the management blade.

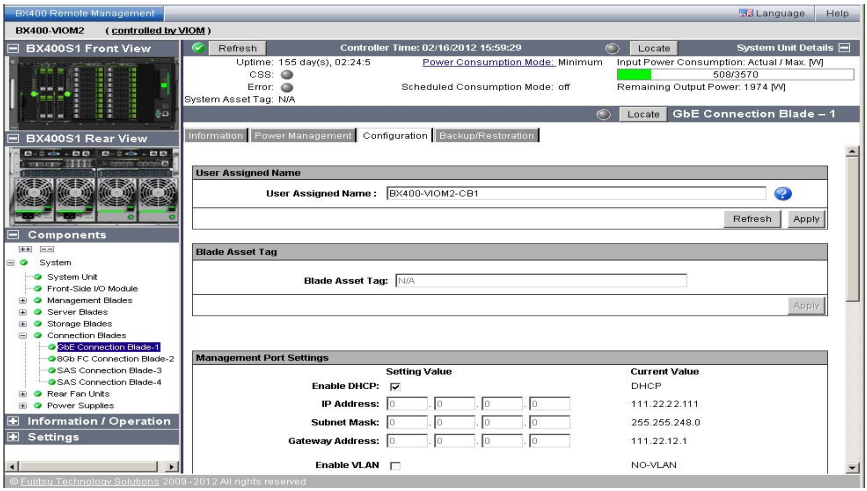

You configure the network parameters of an I/O connection blade as follows:

- Activate the **Enable DHCP** option under **Management Port Settings** if the connection blade is to receive its IP address from a DHCP server. If this option is not activated, you must specify the IP address, subnet mask, and the gateway address.
- Click **Apply** to activate the settings.

4.2.4 Connecting IBP modules

The PY CB Eth Switch/IBP 1Gb 36/8+2, PY CB Eth Switch/IBP 1Gb 36/12 LAN, PY CB Eth Switch/IBP 1Gb 18/6, and PY CB Eth Switch/IBP 10 Gb 18/8 connection blades support two modes:

- "Normal" Layer 2 switch mode
- IBP mode

For information on how to perform the switch, refer to the Switch Blade user manual.

Please note that before you activate blade server management, several uplink ports of a LAN connection blade in IBP mode form a static link aggregation group (LAG). If you activate blade server management using VIOM, all uplink ports of IBP modules will be deactivated and will only be reactivated when networks are defined in the IBP module.

4.2.4.1 Network - Overview

All components such as the management blades and I/O connection blades, the central management station, and the management console must be interconnected via LAN. It is recommended that you set up a management LAN that is separate from the productive LAN, as illustrated in the following figure.

If there are not enough physical ports available, a VLAN-based management LAN can also be implemented.

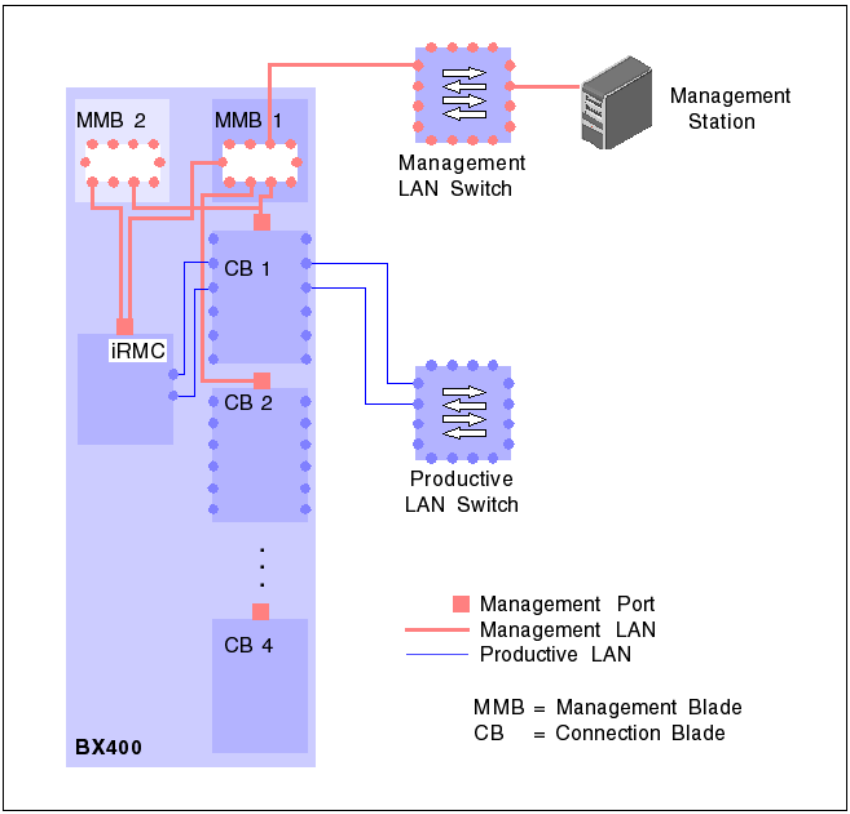

Figure 16: BX400: Network - Overview

The I/O connection blades must be connected to the central management station via their management port.

4.2.5 Switch stacking support

A stacking function is a group set of IBP(s) connected through Infiniband CX4 interface of the HiGig/HiGig+ ports. One of the IBP(s) controlled the operation of the stack modules is called the stack master. Other IBP(s) are belonging to the stack members of the stack group system.

The stacking software configures each device tables and registers to support all switching functions, for example, switching, link aggregation, port

monitoring, spanning tree protocol, VLAN, etc. The entire stack will appear as a single IBP.

You find a detailed description about Switch Stacking in the manual "PRIMERGY BX900 Blade Server Systems- Ethernet Connection Blade Module-IBP version (User`s Guide)", chapter 6.

VIOM supports uplink sets and networks on switch stacks. An uplink set can contain uplink ports from different IBPs in a stack. You can see stacks in the VIOM GUI on the **Setup** tab if you click on a stack member. Then all members of this stack are highlighted: the stack master is yellow shaded; all other members are green shaded. Beside this, **Stack Master** or **Stack Member** is shown in the details area.

To define, modify, or delete uplink sets and networks in a stack, select one stack member and start the desired operation. Details can be found in "Defining network paths (LAN)" on page 277.

4.3 Configurations on the managed BX900 Blade Server

In section "Supported hardware configurations for the connection blades" on page 119, you can find out what combinations of I/O connection blades are supported by VIOM.

To manage a blade server with VIOM, you must carry out the following preparatory steps:

- Configure the management blade or management blades, see section "Configuring the BX900 management blade" on page 121
- Configure the I/O connection blades, see section "Configuring the I/O connection blades" on page 123
- Define the connection of the IBP modules with the network, see section "Connecting IBP modules" on page 126
- Add the blade server to the server list of the ServerView Operations Manager, see section "Adding a server to the ServerView server list" on page 135

4.3.1 Supported hardware configurations for the connection blades

The following connection blades are supported in a BX900:

- The connection blades PY CB Eth Switch/IBP 1Gb 36/8+2, PY CB Eth Switch/IBP 1Gb 36/12, PY CB Eth Switch/IBP 1Gb 18/6, and PY CB Eth Switch/IBP 10 Gb 18/8 as LAN connection blades. LAN connection blades can be inserted in connection bays 1 and 2 of fabric 1 in order to connect the onboard LAN ports and in connection bays 3 and 4 of fabric 2 in order to connect the LAN ports of the first mezzanine card. Irrespective of the defined SMUX setting, the connection blades can be inserted in connection bays 5 and 6 of fabric 3 or in connection bays 7 and 8 of fabric 4 in order to connect the second mezzanine card.

  To define dedicated network connections for server blades with networks, you require LAN connection blades in IBP mode. The PY CB Eth Switch/IBP 1Gb 36/8+2, PY CB Eth Switch/IBP 1Gb 36/12, PY CB Eth Switch/IBP 1Gb 18/6, and PY CB Eth Switch/IBP 10 Gb 18/8 connection blades support a "normal" Layer 2 switch mode and IBP mode.

  If you only intend to work with virtual I/O parameters (Open Fabric mode of the Virtual-IO Manager) and not with dedicated network connections, you can also use the connection blades in switch mode or you can use LAN connection blades that do not support IBP mode. This applies separately to the connection blades in each fabric. In other words, you can operate connection blades in IBP mode in fabric 1 and connection blades in switch mode in fabric 2.
- You must install an 8 GB Brocade switch blade in fabric 2 and/or 3 in order to support a Fibre Channel mezzanine card.

4.3.1.1 LAN hardware configuration

Using IBP modules

All fabrics have identical configuration restrictions in relation to the IBP modules. The following table shows you which configurations are supported in

fabric 1, fabric 2, fabric 3, and fabric 4.

Bay 1 (Bay 3) (Bay 5) (Bay 7)	Bay 2 (Bay 4) (Bay 6) (Bay 8)
PY CB Eth Switch/IBP 1Gb 36/8+2 in IBP mode	PY CB Eth Switch/IBP 1Gb 36/8+2 in IBP mode (optional[1])
PY CB Eth Switch/IBP 1Gb 36/12 in IBP mode	PY CB Eth Switch/IBP 1Gb 36/12 in IBP mode (optional[1])
PY CB Eth Switch/IBP 1Gb 18/6 in IBP mode	PY CB Eth Switch/IBP 1Gb 18/6 in IBP mode (optional[1])
PY CB Eth Switch/IBP 10 Gb18/8 in IBP mode	PY CB Eth Switch/IBP 10 Gb 18/8 in IBP mode (optional[1])

[1] Even though this constellation is possible, it does not make sense because only the first port can be used from the IBP module in bay 1, 3, 5 or 7.

Operation in Open Fabric mode

Open Fabric mode refers to the use of non-VIOM-capable LAN modules together with the Virtual-IO Manager. In this case, you can only work with virtual I/O parameters, and not with dedicated network connections.

Therefore, if you do not want to use any routing, all combinations are possible even with non-VIOM-capable I/O connection blades. You must take the following into account:

- You cannot mix LAN and IBP modules within a fabric.
- You can only assign profiles if no networks are defined for their ports.

  You can force the assignment of a profile with network definitions by selecting the **Ignore ext. LAN connections** option in the profile selection dialog box or by answering the subsequent query accordingly. In doing so, however, the network definitions are ignored.

- You can operate LAN models of different sizes within a fabric, although you cannot use all download links from the larger model.

4.3.1.2 Fibre Channel hardware configuration

You can install the Fibre Channel switch blades in fabric 2 or fabric 3. They can be operated in two different modes:

- in normal switch mode
- in Access Gateway mode

As a result, different combination options are possible. Fabric 2 and fabric 3 have identical configuration restrictions in relation to the Fibre Channel switch blades. The following table shows you which configurations are supported in fabric 2 and fabric 3:

Bay 3 (Bay 5)	Bay 4 (Bay 6)
Access Gateway mode	Access Gateway mode (optional[1])
Standard switch mode	Standard switch mode (optional[1])

[1]Even though this constellation is possible, it does not make sense because only the first port can be used from the Fibre Channel switch blade in bay 3 or 5.

4.3.2 Configuring the BX900 management blade

The blade server must be fitted with at least one S1 management blade.

The management blade has two user-friendly user interfaces: a Web interface and a Remote Manager interface per Telnet or SSH protocol (a connection to the management blade is established per Telnet or SSH). You can find a detailed description on the interfaces in the ServerView Management Blade manual for BX900.

Check the following settings in the management blade:

- The management blade must be installed with a specific firmware version. To find out which firmware version you need, see the release notes included.You can check the firmware version in the Remote Manager of the management blade. To do this, choose **(1) Management Agent – (2) Management Blade**. The **Management Blade Firmware Version** parameter displays the installed firmware version. You can also check

the version via the Web interface.

If your firmware version is lower than the version required, you must update it to the required firmware version before you activate VIOM management for the blade server. You can run an update via the Remote Manager of the management blade. For a MMB S1, choose **(4) TFTP update**. First specify the IP address of the TFTP server by selecting option **(1) TFTP Server IP Adress**. Then specify path and name of the file with the new firmware version by selecting option **(2) Management Blade Image File Name**. After this select option **(3) Management Blade Update Enable**.

- You must configure the management blade and connect it to the local network so that the management LAN port of the management blade can be accessed from the central management station on which the ServerView Operations Manager and VIOM is installed.
- Access to the management blade via Telnet or SSH is essential for VIOM. You must configure the management blade via its Web interface. To do this, choose **Settings – System Unit – Network Interface – Management LAN**. On the **Ethernet** tab, check the telnet and SSH settings in the **Text based Access** section.
- Ensure that the **Automatic Inventory Retrieval** parameter is set to **automatic**, so that the management blade automatically generates the inventory information required by VIOM when the server blade is inserted.

  

  VIOM needs this inventory information specifically for the server blades that support the I/O virtualization.

  To check the parameter setting in the Remote Manager of the management blade, choose **(1) Management Agent – (3) System Information – (10) Automatic Inventory Retrieval**.

  To check the parameter setting in the Web interface open the **Components** section and select **System– System Unit**. The **System Information** section on the right contains the option **Automatic Inventory Retrieval**. If the value of this is not set to **automatic**, please set it to

automatic and remove and insert all blades or perform a power-off and power-on of the blade server chassis.

- To ensure that the blade server names in the ServerView server list are unique, you must assign a system name to the blade server chassis. System names are assigned when the management blade is configured.

4.3.3 Configuring the I/O connection blades

The connection blades in fabric 1 and the optional connection blades in fabric 2, 3, and 4 must be configured and connected to the local network so that they can be accessed from the central management station on which the ServerView Operations Manager and VIOM are installed.

You can configure the network parameters of an I/O connection blade using the easy-to-use Web interface or you can use the Remote Manager of the management blade.

Configuration using the Web interface

It is important that the I/O connection blades of the BX900 can be accessed from the management station via their management port.

You can configure the network parameters of the management port of an I/O connection blade using the Web interface of the management blade.

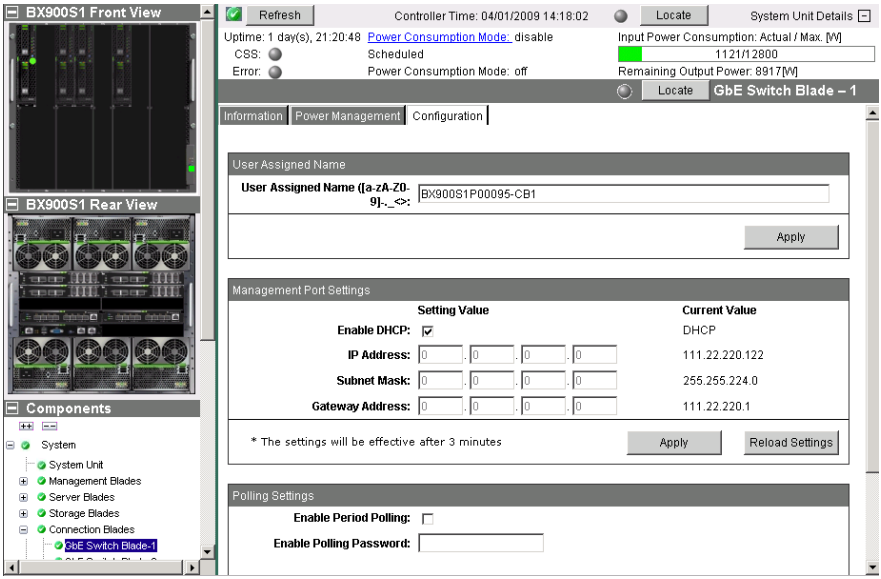

You configure the network parameters of an I/O connection blade as follows:

- Activate the **Enable DHCP** option under **Management Port Settings** if the connection blade is to receive its IP address from a DHCP server. If this option is not activated, you must specify the IP address, subnet mask, and the gateway address.
- Click **Apply** to activate the settings.

Configuration using the Remote Manager

You can also configure the network parameters of the I/O connection blades using the Remote Manager of the management blade.

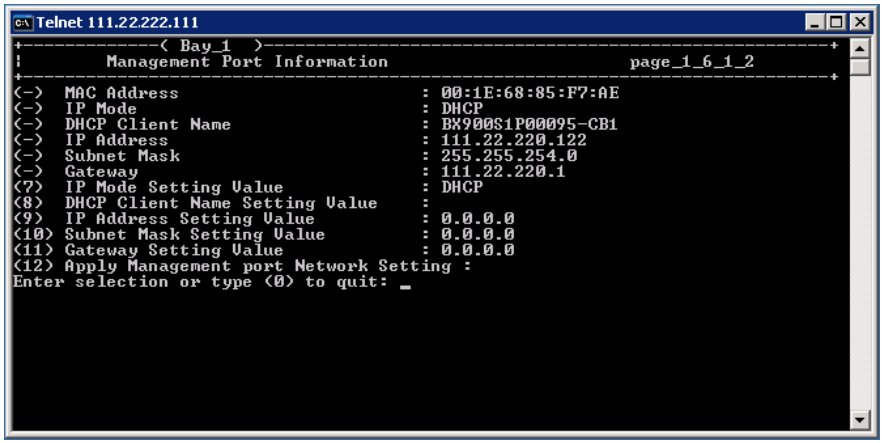

You configure the network parameters of an I/O connection blade as follows:

- In the Remote Manager, choose **(1) Management Agent – (6) Connection Blade**.
- Select an I/O connection blade.
- Choose **(2) Management Port Information**.
- Make the relevant entries for the following parameters
  - **IP Mode Setting Value**
  - **IP Address Setting Value**
  - **Subnet Mask Setting Value**
  - **Gateway Setting Value**
- Activate the settings by clicking **Apply Management port Network Setting**.

  For Fibre Channel switch modules (select **FC Switch Blade**), you must also specify access data for the parameters **Set Login Username** and **Set Login Password** beforehand for the corresponding I/O connection blade in the management blade.

4.3.4 Connecting IBP modules

The PY CB Eth Switch/IBP 1Gb 36/8+2, PY CB Eth Switch/IBP 1Gb 36/12 LAN, PY CB Eth Switch/IBP 1Gb 18/6, and PY CB Eth Switch/IBP 10 Gb 18/8 connection blades support two modes:

- "Normal" Layer 2 switch mode
- IBP mode

For information on how to perform the switch, refer to the Switch Blade user manual.

Please note that before you activate blade server management, several uplink ports of a LAN connection blade in IBP mode form a static link aggregation group (LAG). If you activate blade server management using VIOM, all uplink ports of IBP modules will be deactivated and will only be reactivated when networks are defined in the IBP module.

4.3.4.1 Network - Overview

All components such as the management blades and I/O connection blades, the central management station, and the management console must be interconnected via LAN. It is recommended that you set up a management LAN that is separate from the productive LAN, as illustrated in the following figure.

If there are not enough physical ports available, a VLAN-based management LAN can also be implemented.

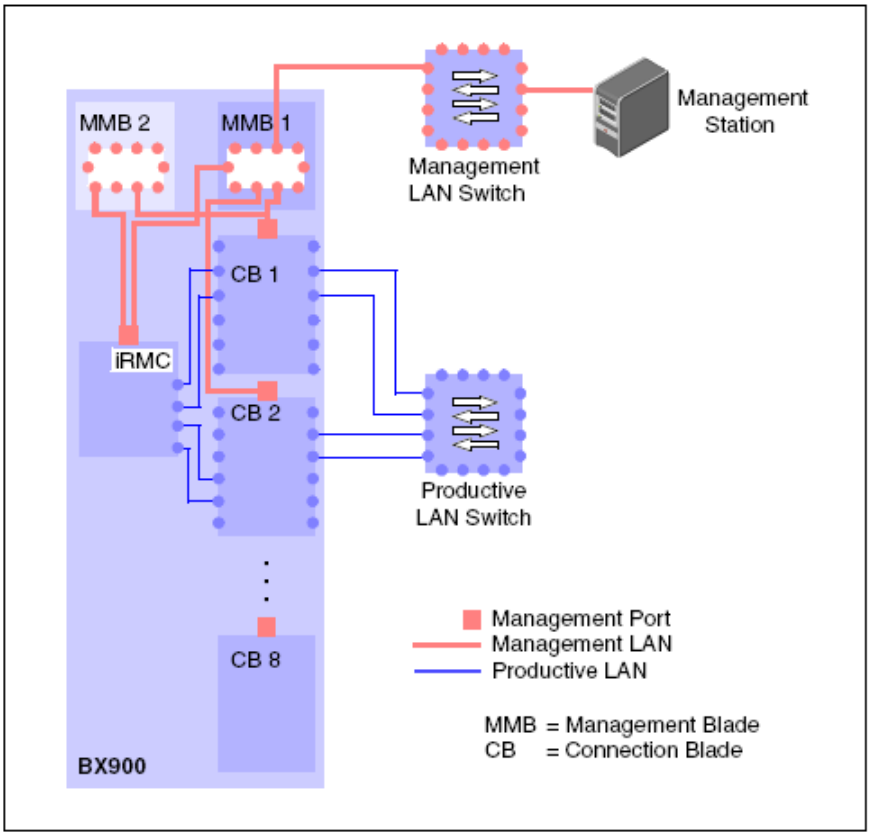

Figure 17: BX900: Network - Overview

The I/O connection blades must be connected to the central management station via their management port.

4.3.5 Switch stacking support

A stacking function is a group set of IBP(s) connected through Infiniband CX4 interface of the HiGig/HiGig+ ports. One of the IBP(s) controlled the operation of the stack modules is called the stack master. Other IBP(s) are belonging to the stack members of the stack group system.

The stacking software configures each device tables and registers to support all switching functions, for example, switching, link aggregation, port monitoring, spanning tree protocol, VLAN, etc. The entire stack will appear as a single IBP.

You find a detailed description about switch stacking in the "PRIMERGY BX900 Blade Server Systems- Ethernet Connection Blade Module- IBP version (User`s Guide)", chapter 6.

VIOM supports uplink sets and networks on switch stacks. An uplink set can contain uplink ports from different IBPs in a stack. You can see stacks in the VIOM GUI on the **Setup** tab if you click on a stack member. Then all members of this stack are highlighted: the stack master is yellow shaded; all other members are green shaded. Beside this, **Stack Master** or **Stack Member** is shown in the details area.

To define, modify, or delete uplink sets and networks in a stack, select one stack member and start the desired operation. Details can be found in "Defining network paths (LAN)" on page 277.

Using stacked SB11 Connection Blades in IBP Mode

VIOM also supports stacked SB11 connection blades in IBP mode, but the configured stack must not span over several chassis. (Stacked SB11 connection blades in switch mode are also supported since VIOM handles them transparently. In this case the stack can span over several chassis.)

For correct handling by VIOM the SNMP agent on the stacked SB11 connection blades must be running and the SNMP community for the SNMP agent on the stack must be configured to be identical to the SNMP community of the blade server management blade. If this is not done, ServerView Operations Manager cannot retrieve the necessary information for VIOM in order to detect the stack correctly and the stack members might be handled as single connection blades which will result in errors.

It is important to configure the SB11 IBP stack before managing the chassis by VIOM. Also the topology of the stack (which connection blades belong to the stack) should not be changed, while the chassis is managed. If the stack gets broken while the chassis is managed by VIOM, this is handled as a stacking error by VIOM. In this case the stack needs to be repaired.

It is also important that the members of a stack have a unique "user assigned" name in the MMB and that the stack member IDs of the participating connection blades are set to 1 before building the stack.

The typical procedure for correct handling by VIOM is:

- Configure the stack by enabling the blade server internal stacking links and by connecting SB11 connection blades in different fabrics of the same chassis by special stacking cables.
- Check that the blade server management blade correctly shows the status of the stack. For this it is important that all members of the stack show the same IP address, which must be the IP address of the stack master. The stacking status in MMB must be **Master** for exactly one member of the stack and must be **Member** for all other members of the stack. As long as these two conditions are not fulfilled, VIOM will not be able to handle the stack correctly.
- Update the information for this blade server chassis in ServerView Operations Manager by performing the “explore” action for the blade server chassis
- Perform the “Refresh” action in VIOM user interface so that also VIOM has the updated information about the changed configuration. The last two actions will also be done automatically when waiting about two or three minutes (this also depends on the refresh interval configured in ServerView Operations Manager).

It is not supported to change the stack configuration while the blade server chassis is managed by VIOM. If it is needed to change the stack configuration, the chassis is to be unmanaged in VIOM before changing the configuration.

4.4 Configurations on the managed PRIMERGY rack server

A PRIMERGY rack server must fulfill the following requirements in order to be manageable by VIOM:

- The PRIMERGY rack server must be in the list of supported server types.
- The system BIOS and the iRMC firmware must fulfill VIOM requirements. Details concerning the required BIOS version and iRMC firmware version might be found in the release notes.
- The iRMC of the PRIMERGY rack server must be reachable by LAN from the central management station where VIOM is installed. VIOM uses the RMCP protocol (IPMI over LAN) to communicate with the iRMC of the PRIMERGY rack server. If any routers or firewalls are involved, this protocol must be allowed between the VIOM central management station and the managed PRIMERGY rack server.

  By default the iRMC is configured to receive an IP address via DHCP. If no DHCP server is running you must configure a static IP address, netmask and default gateway manually for the iRMC of every PRIMERGY rack server that is to be managed by VIOM.
- The PRIMERGY rack server and especially the iRMC must be added to the server list of ServerView Operations Manager. ServerView Operations Manager must be configured in such a way that it can communicate with the iRMC. This means that ServerView Operations Manager must know a valid user account and password that allow RMCP communication with the iRMC of that server.

  ServerView Operations Manager must show the PRIMERGY rack server as manageable by VIOM.

  A VIOM icon indicates whether this PRIMERGY rack server is manageable by VIOM or not:

	This PRIMERGY rack server is managed by VIOM.
	This PRIMERGY rack server is manageable by VIOM but management is inactive.
No icon	This PRIMERGY rack server could not be managed by VIOM or its state could not be provided.

For further information on the server list, see the "ServerView Operations Manager" user guide.

- For correct operation in case of power loss (AC failure) for PRIMERGY rack servers, the iRMC of these servers must be able to send SNMP traps to the central ServerView management station with ServerView Operations Manager and VIOM running on it. If any routers or firewalls are involved, these must be configured in such a way that SNMP traps can be delivered to the ServerView management.

  If a PRIMERGY rack server is managed by VIOM, as well as the user account used for contacting the iRMC, you must specify the IP address or host name of the ServerView management station. The supplied address is registered on the iRMC as a trap destination. In addition, on each iRMC the SNMP community for SNMP traps must match the configuration of ServerView Operations Manager.

  It has to be ensured that SNMP traps sent by the iRMC are allowed to pass all routers and fire walls to reach the local LAN port to which the specified IP address is associated.

- If an operating system is installed and running, ServerView Operations Manager should have SNMP access to the server. It is preferable to install ServerView agents on the operating system.

When managing multiple PRIMERGY rack servers via VIOM, the virtualization of the I/O ports together with the definition of the boot device configuration (Fibre Channel boot, iSCSI boot, PXE boot) by VIOM allows easy shifting of an application (OS boot image) from one server to another. To make this possible, the servers that should be able to run the same image should have identical or nearly identical hardware configuration regarding I/O controllers (Fibre Channel controllers, additional LAN controllers, Converged Network Controllers). If, for example, your VIOM server profile uses two onboard LAN ports and two Fibre Channel ports on PCI card 2, all PRIMERGY rack servers that should be able to take these server profiles must have at least two onboard LAN ports and a Fibre Channel controller in the PCI slot 2. Additional controllers in other PCI slots do not cause problems when assigning a VIOM server profile. Missing I/O ports (I/O ports that are configured in the VIOM server profile for which there are no hardware ports) will make it impossible to successfully assign the profile.

Although VIOM does not guarantee that boot images created for one server type also work with another server type, VIOM also supports the assignment of server profiles for blades on PRIMERGY rack servers and vice versa. For further information, see the section "VIOM server profile mapping" on page 132.

4.5 VIOM server profile mapping

Although server profiles for PRIMERGY rack server models and blade servers differ with regard to I/O channel types, profiles for PRIMERGY rack servers can be assigned to blade servers and profiles for blade servers can be assigned to PRIMERGY rack servers if the target server provides the same I/O controllers and meets certain requirements.

To determine if such an assign is possible, it is important to know how I/O channels of blade server mezzanine cards are mapped to I/O channels of PCI add-on cards and vice versa. Onboard I/O channels are taken as they are. Of course, the device type of the destination ports have to be compatible to the source ports. The used I/O channel mapping is shown in the following table:

Mezzanine card in blade server	PCI controller in PRIMERGY rack server
Mezzanine card 1, I/O port 1	PCI controller 1, I/O port 1
Mezzanine card 1, I/O port 2	PCI controller 1, I/O port 2
Mezzanine card 1, I/O port 3	PCI controller 1, I/O port 3
Mezzanine card 1, I/O port 4	PCI controller 1, I/O port 4
Mezzanine card 2, I/O port 1	PCI controller 2, I/O port 1
Mezzanine card 2, I/O port 2	PCI controller 2, I/O port 2
Mezzanine card 2, I/O port 3	PCI controller 2, I/O port 3
Mezzanine card 2, I/O port 4	PCI controller 2, I/O port 4

Figure 18: Mapping of IO-channels

- Profiles with onboard CNA configuration cannot be assigned to servers which do not have an onboard CNA controller.
- Only I/O channels on master slots are mapped. So if a server profile for a BX960 S1 contains I/O channels for the slave slot they are ignored when the profile is assigned to a PRIMERGY rack server.

4.6 PCI slot location in PRIMERGY rack servers

The PCI slots of the supported PRIMERGY rack servers are located as follows.

RX200 S7

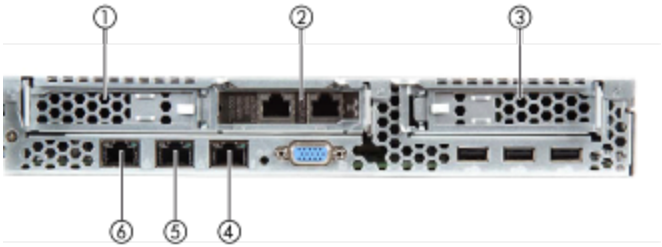

Figure 19: RX200 S7 connection panel on the rear

1	PCI slot 3
2	PCI slot 2
3	PCI slot 1
4	Management LAN connector (iRMC)
5	Standard LAN connector (LAN 2)
6	Shared LAN connector (LAN 1)

RX300 S7

Figure 20: RX300 S7 connection panel on the rear

1	PCI slot 1
2	PCI slot 2
3	PCI slot 3
4	PCI slot 4
5	Shared LAN connector (LAN 1)
6	Standard LAN connector (LAN 2)
7	Management LAN connector (iRMC)
8	PCI slot 5
9	PCI slot 6

TX300 S7 and RX350 S7

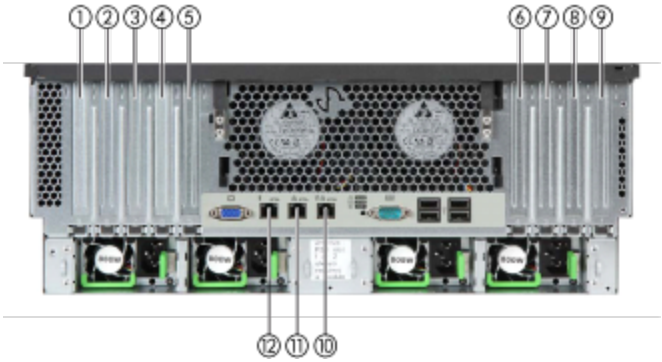

Figure 21: TX300 S7 connection panel on the rear

1	PCI slot 9
2	PCI slot 8
3	PCI slot 7
4	PCI slot 6
5	PCI slot 5
6	PCI slot 4
7	PCI slot 3
8	PCI slot 2
9	PCI slot 1
10	Shared LAN connector (LAN 1)
11	Standard LAN connector (LAN 2)
12	Management LAN connector (iRMC)

4.7 Adding a server to the ServerView server list

In order to manage a server with VIOM, you have to add it to the server list of the ServerView Operations Manager after you have configured it.

You add the server to the ServerView server list using the server browser of the ServerView Operations Manager. To do this, choose **Administration – ServerBrowser** to open the **Server Browser** properties window.

In this window, you have to specify the IP address and the system name of the management blade.

You can find a detailed description on the server browser in the ServerView Operations Manager user manual.

If the server is added to the ServerView server list, then this server is also automatically added to the VIOM-specific server group **VIOM Manageable**.

In order to manage a server, it is essential to move it from the **VIOM Manageable** server group to the second VIOM-specific server group **VIOM Managed**. In the context of VIOM, we refer to this as VIOM management activation. More information on activating and deactivating VIOM management for a server is described in chapter "Managing servers with VIOM" on page 257.

5 Virtual-IO Manager user interface

This chapter contains a general introduction to the Virtual-IO Manager (VIOM) describing:

- The structure of the Virtual-IO Manager start page
- The tabs, wizards, dialog boxes, buttons, and context menus that are available when working with VIOM
- The meaning of the various icons

5.1 Virtual-IO Manager main window

Once you have launched the Virtual-IO Manager, the following main window is displayed:

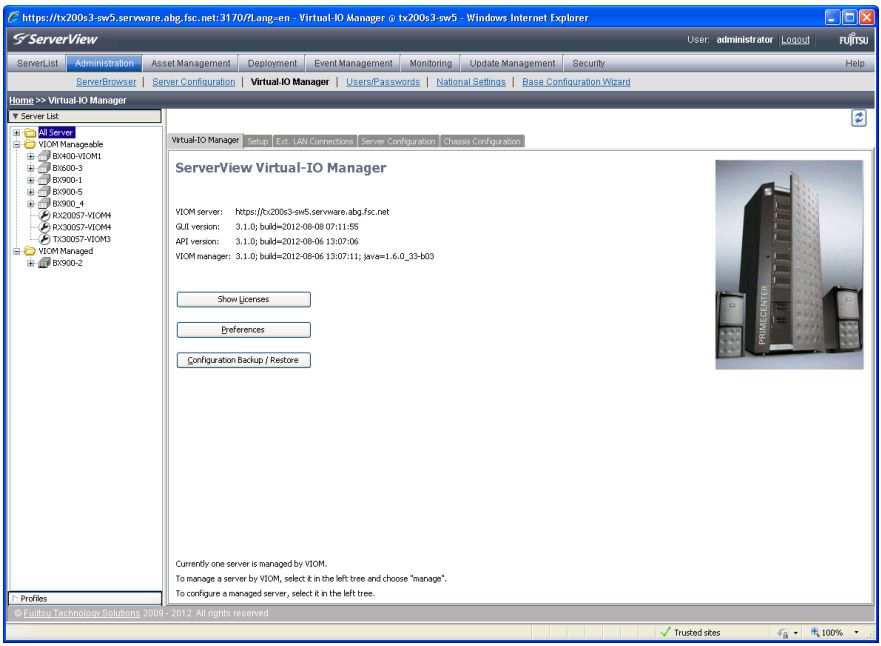

Figure 22: VIOM main window

The window is divided into three areas:

- The ServerView Suite header
- The menu bar
- The work area with the tree structure on the left and tabs and views on the right

The menu bar below the headline allows you to navigate between the various functions of the ServerView Operations Manager:

- **ServerList**
- **Administration**
- **Asset Management**
- **Event Management**
- **Monitoring**
- **Update Management**

Depending on which menu you select, the individual menu items for this menu are displayed in the row below the menu bar.

For more information on the menus of the menu bar, see the "ServerView Operations Manager user manual".

Depending on what you select on the left in the work area, the following is displayed in a tree structure:

- server list
- server profiles

You switch between the two views by clicking the **Server List** and **Profiles** buttons.

The area on the right displays the following depending on what is selected in the area on the left:

Button	Selection on the left	Display on the right
Server List		Five tabs are displayed in the area on the right (see "Tabs" on page 142).
	All Server server group	Only the **Virtual-IO Manager** tab is activated, which contains general information on VIOM.
	VIOM Manageable or **VIOM Managed** server group	Same as if all servers of the group were selected.
	Object(s) in the **VIOM Manageable** server group	Only the **Virtual-IO Manager** and **Setup** tabs are activated in the area on the right. You can activate VIOM management of objects on the **Setup** tab.
	Object(s) in the **VIOM Managed** server group	The **Virtual-IO Manager**, **Setup**, and **Server Configuration** tabs are activated in the area on the right. The **Ext. LAN Connections** and **Chassis Configuration** tabs are only enabled if a single blade server is selected.
Profile	**Profiles** group	An overview of the previously defined profiles
	A profile	Only the corresponding profile

It is not possible to select unmanaged and managed servers in the server list at the same time.

5.2 Tree view

The tree view is on the left of the work area in the Virtual-IO Manager and provides various views.

- **Server list** (default)
- **Profiles**

You switch between the two views by clicking the **Server List** and **Profiles** buttons. The **Server List** button is always above the tree structure. The **Profiles** button is:

- Below the tree structure if the server list is displayed in this area
- Directly above the tree structure if the profiles are displayed in this area.

The display in the area on the right changes according to what you select in this area.

The icons in front of the objects in the tree structure match the icons in the ServerView server list. You can find a table describing the icons and their meaning in the ServerView Operations Manager user manual.

5.2.1 Tree structure (Server List)

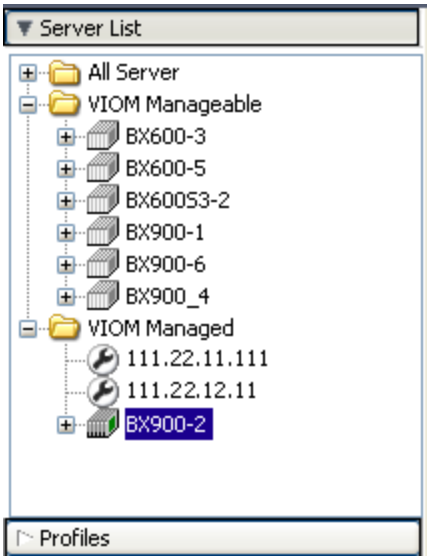

Figure 23: Server List (tree structure)

When you select **Server List** (standard) on the left of the work area, the **All Server** server group is displayed according to the ServerView server list as well as the two VIOM-specific server groups **VIOM Manageable** and **VIOM Managed**.

All Server	Includes the servers according to the ServerView server list
VIOM Manageable	Includes the servers that can be managed using VIOM, but are currently not managed
VIOM Managed	Includes the servers that are managed using VIOM. This server group contains the servers that were selected in the **VIOM Manageable** group for which the **Manage** action was successfully carried out. The servers managed by VIOM are then no longer listed in the **VIOM Manageable** group.

5.2.2 Tree structure (Profiles)

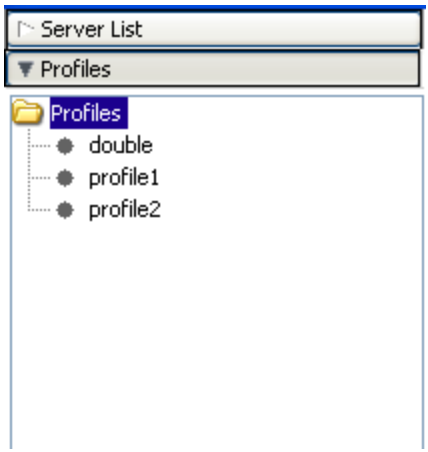

Figure 24: Profiles (tree structure)

When you select **Profiles** on the left of the work area, the defined server profiles are displayed. The folder is empty as long as a server profile has not yet been created.

- Select a folder to show all contained profiles in the profiles table in the **Server Profiles** view in the area on the right of the **ServerView Virtual-IO Manager** window. By using the context menu in the area on the left, you can create a new profile.
- Select a profile to show it in the profiles table. By using the context menu, you can edit and delete this profile or make a copy of it.

5.3 Tabs

If **Server List** is activated in the tree view, Virtual-IO Manager provides the following tabs:

- **Virtual-IO Manager** tab
- **Setup** tab
- **Ext. LAN Connections** tab with:
  - **Graphic** tab
  - **Details** tab
- **Server Configuration** tab
- **Chassis Configuration** tab

Initially the **Virtual-IO Manager** tab is activated.

5.3.1 Virtual-IO Manager tab

The **Virtual-IO Manager** tab contains general information on VIOM. It is the only tab that is always enabled if **Server List** is selected in the tree view.

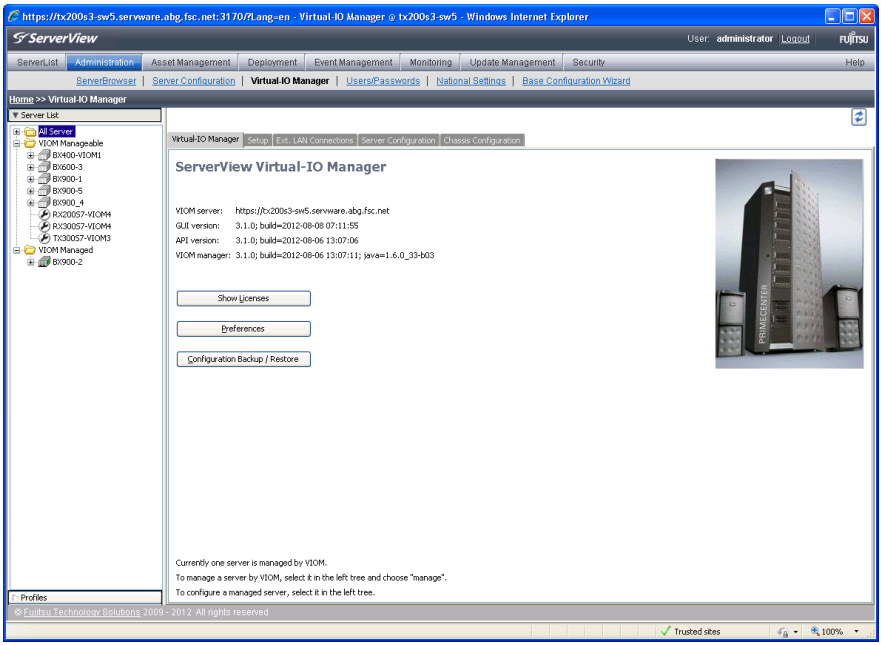

Figure 25: Virtual-IO Manager tab

Show Licenses

Displaying information on the licenses assigned (see section "Displaying license information" on page 275)

Preferences

Showing and modifying user preferences (see section "Preferences dialog box" on page 238)

Configuration Backup / Restore

Saving and restoring the blade server configuration and the server profiles (see chapter "Saving and restoring" on page 303)

5.3.2 Setup tab

Using this tab, you specify whether an object is to be managed with VIOM or not.

- For a blade server or PRIMERGY rack server in the **VIOM Manageable** server group, you can activate the administration of VIOM using the **Setup** tab by adding this server to the **VIOM Managed** server group.
- For a blade server or PRIMERGY rack server in the **VIOM Managed** server group, you can deactivate management with VIOM by adding this server to the **VIOM Manageable** server group.

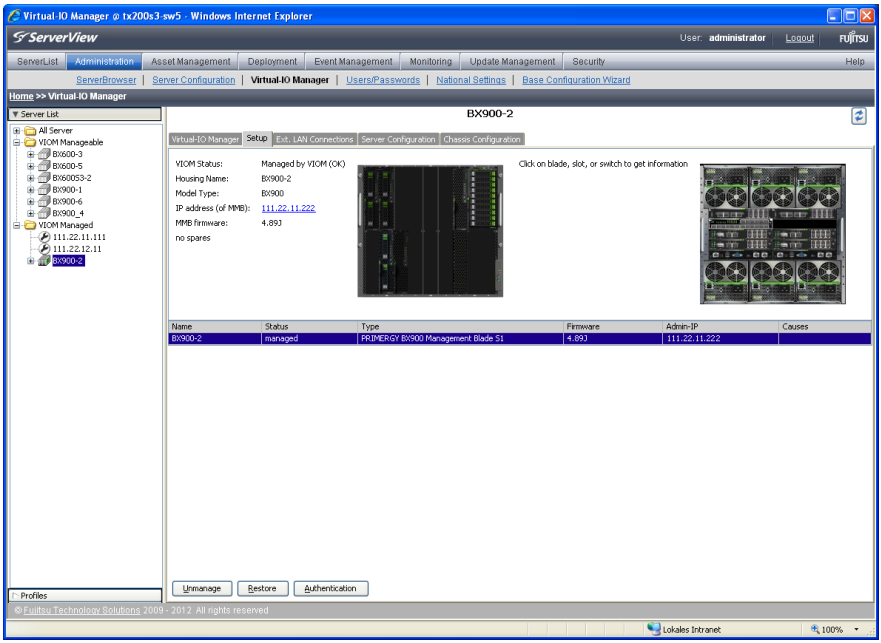

Figure 26: Setup tab

The **Setup** tab contains information on the servers selected in **Server List**.

The table contains one entry for each of the blade or PRIMERGY rack servers selected in the server list. For selected server blades, the cor-

responding blade server is shown. The table provides the following information:

Column	Significance
Name	Name of the server
Status	Manage status of the server
Type	Product name
Firmware	Firmware version
Admin-IP	IP address to administrate the server
Causes	Causes. For the possible values and its meaning see the table below

If a single server is selected in the table above the table, there are more details about this server available:

- The front and rear (rear only for blade server) according to the actual configuration
- Information on the model
- The IP address and firmware version of the management blade or PRIMERGY rack server
- The number of spare slots defined (for blade servers only)
- State causes (for PRIMERGY rack servers only)

If a single blade server is selected in the table, slots can be selected in the graphic to get more information about the slot, the server blade, or switch blade which is plugged into this slot:

- Number of slot or connection bay
- Stacking info (if available)
- Product name if selected slot is not empty
- IP address (if available)
- BIOS version for server blades or Firmware version for switches
- Manage State

- Fault State
- State Causes

The (**Causes** or **State Causes**) state cause is not at all a fault. It is simply a state marker. The following values are defined:

Cause	Significance
data base doesn't match hardware	Configuration of switch/server does not match with database.
not supported	Switch/server model is not supported.
model does not match slot configuration	The model is supported, but some other requirements are not met.
slot is empty	No hardware plugged, although a configuration exists.
profile assigned	Switch/server is used by an assigned server profile.
no configuration	No configuration of switch/server is stored inside the database.
no ServerView entry	No entry of blade server inside of ServerView Operations Manager.
different switches in fabric	The switch is incompatible with another switch in the same fabric.
inapplicable profile assigned	The assigned server profile is incompatible with the plugged blade server.
stacking error	The switch configuration is not compatible with the stacking groups.

Depending on what you select, the following buttons are available, that allow you to carry out various actions.

Manage

Activates the management of the selected servers. The **Authentication** dialog box opens. In this dialog box, you enter the user names and

passwords, which VIOM can use to access these servers. For details see "Authentication dialog box (single blade server)" on page 231, "Authentication dialog box (PRIMERGY rack server)" on page 234, and "Authentication dialog box (PRIMERGY rack server and blade server)" on page 236. The servers are added to the **VIOM Managed** server group.

Unmanage

Deactivate management with VIOM by adding the selected servers to the **VIOM Manageable** server group.

Authentication

Opens the **Authentication** dialog box. In the **Authentication** dialog box, you enter the data for authentication (e.g. user names, passwords, ports), which VIOM can use to access the selected servers. For details see "Authentication dialog box (single blade server)" on page 231, "Authentication dialog box (PRIMERGY rack server)" on page 234, and "Authentication dialog box (PRIMERGY rack server and blade server)" on page 236.

Restore

Use the **Restore** button to write the configuration and virtualization data (virtual I/O addresses and possibly also boot parameters) saved internally in VIOM database to all IBP modules and slots of the selected blade server. For the selected PRIMERGY rack servers, the virtualization data are rewritten.

Restore Slot

blade server only (front-side slot selected)

If you select a slot, you can use the **Restore Slot** button to rewrite the virtualization data (virtual I/O addresses and possibly also boot parameters) of the assigned profile if a problem occurs.

Restore IBP

blade server only (IBP selected)

If you select an I/O connection blade and this module does not have the **wrong model** or **stacking error** state in **State Causes**, you can use the

Restore IBP button to write the configuration saved internally in the VIOM database to the IBP module.

Delete Configuration

blade server only (switch or switch slot selected)

If you select a slot for an I/O connection blade and the slot has the **empty slot**, **wrong model**, or **stacking error** state, you can use the **Delete Configuration** button to delete the configuration saved internally in the VIOM database.

Video Redirection

PRIMERGY rack server and server blades

Opens a new window where the console output of the server is shown.

5.3.3 Ext. LAN Connections tab

On the **Ext. LAN Connections** tab, you configure the LAN modules with the external devices. Two tabs are available for doing this. It is only enabled if a single blade server is selected in the server list.

- **Graphic** tab

  This tab schematically displays the connection blades (IBPs) and ports. Using this tab, you can create new network definitions, modify or delete definitions as well as copy definitions from one connection blade to another one.

- **Details** tab

  This tab contains a table with the present definitions. Using this tab, you can create new network definitions, or modify and delete existing ones.

Also on this tab, you find the **Copy Configuration** button for copying all networks of a chassis. For more information see section "Copying configuration" on page 288

5.3.3.1 Graphic tab on Ext. LAN Connections tab

The **Graphic** tab contains a schematic display of the I/O connection blades and ports. Using this tab, you can create new network definitions, edit or

delete existing ones, as well as copy definitions from one connection blade to another.

In the upper left part of the **Graphic** tab there is the rear view of the selected chassis. If an IBP is selected in this graphic, its uplink sets and networks are shown in the table on the right. Additional information about the selected IBP is shown on the left of the graphic. At the bottom of the tab there is a schematic view of the uplink ports of the selected IBP.

If the selected IBP belongs to a switch stack, all members of this stack are highlighted by coloring them (yellow for the stack master, green for other stack members). In the port area, the uplink ports of all stack members are displayed.

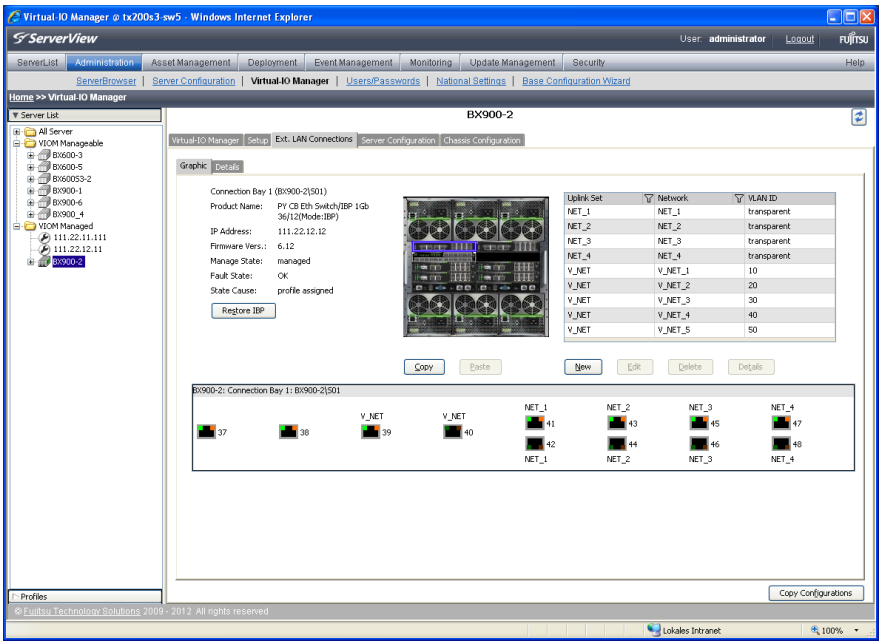

Figure 27: Graphic tab on the Ext. LAN Connections tab

New

Click **New** to define a new uplink set. The **Create network** wizard is launched.

Edit

Click **Edit** to edit an existing uplink set. The **Edit Uplink Set** wizard is launched in which you can change the configuration.

Delete

Click **Delete** to delete networks or uplink sets.

Details

Click **Details** to get more information on the selected uplink set. The **Uplink Set <name of uplink set>** window opens.

For more information on this, see chapter "Defining network paths (LAN)" on page 277.

5.3.3.2 Details tab on Ext. LAN Connections tab

The **Details** tab contains a table with the existing definitions. Using this tab, you can also define new network definitions, or modify and delete existing ones.

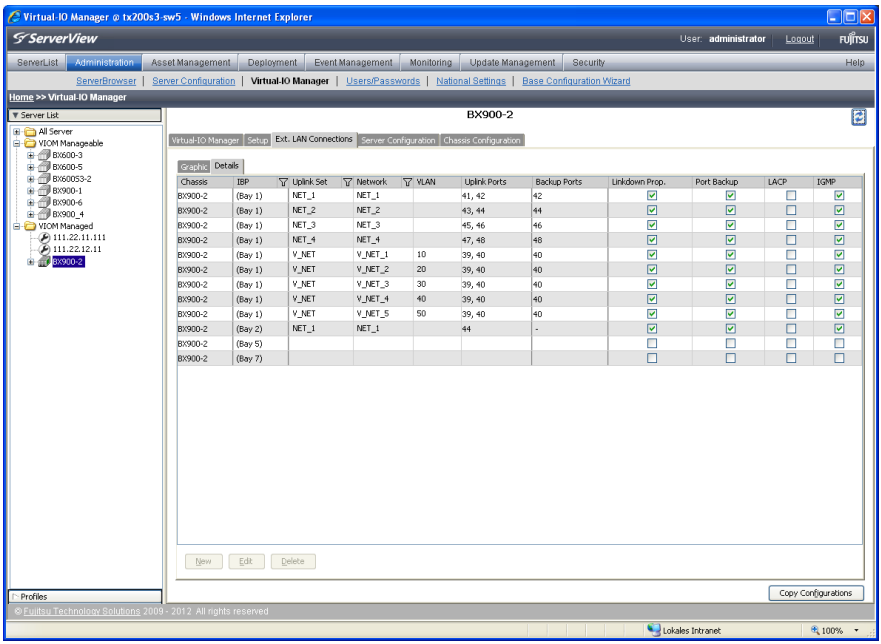

Figure 28: Details tab on the Ext. LAN Connections tab

You can sort the table by the columns **IBP**, **Uplink Set**, and **Network**.

New

Click **New** to define a new uplink set. The **Create network** wizard is launched.

Edit

Click **Edit** to edit an existing uplink set. The **Edit Uplink Set** wizard is launched in which you can change the configuration.

Delete

Click **Delete** to delete networks or uplink sets.

For more information on this, see chapter "Defining network paths (LAN)" on page 277.

5.3.4 Server Configuration tab

On the **Server Configuration** tab, you can do the following:

- Assign server profiles to the slots of a blade server or to a PRIMERGY rack server
- Delete server profile assignments
- Define a new server profile directly from here
- Define spare slots on blade servers
- Move a server profile from one slot to a suitable spare slot (start a server profile failover), only for blade server
- Switch server blades or PRIMERGY rack servers on and off

For more information on this, see chapter "Defining and assigning server profiles" on page 289.

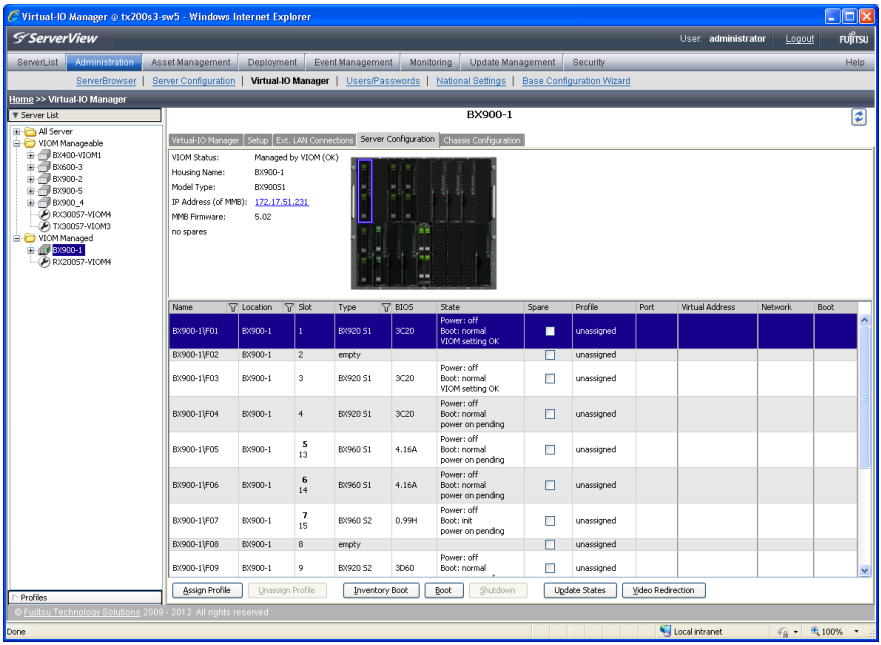

Figure 29: Server Configuration tab

The table contains one entry for each slot of the blade servers selected in the server list and one entry for each PRIMERGY rack server selected in the server list. If only slots of one blade server or only one PRIMERGY rack server is selected in the table, more information about this server is shown above the table. The table provides the following information:

Column	Significance
Name	Name of the server blade or PRIMERGY rack server
Location	For server blades, the name of the chassis containing the server blade; otherwise empty
Slot	Slot number; blank for PRIMERGY rack servers
Type	Type of the server blade or PRIMERGY rack server or blank if the slot is empty
BIOS	BIOS version of the server blade or PRIMERGY rack server

Column	Significance
State	Power state
Spare	Indicates whether or not the slot is a spare slot. You change the definition of a slot as a spare slot directly in the table (see section "Moving tasks using the server profile failover" on page 332).
Profile	Name of the assigned server profile, or unassigned if a profile is not yet assigned to the slot or PRIMERGY rack server
Port	If a profile is assigned to the slot or PRIMERGY rack server, the column contains an entry for each port with the port number.
Virtual Address	If a profile is assigned to the slot or PRIMERGY rack server, this column contains an entry with the virtual MAC address or the virtual WWN address for each port.
Network	If a profile is assigned to the slot or PRIMERGY rack server, the column contains an entry for each port with the names of the networks that use this port. If this port is also used by service networks, those names are enclosed in **S(** and **)**. Tagged VLAN network names are enclosed in **T(** and **)**. Example: **network,S(service),T(tagged)**

Buttons

Assign Profile

Click **Assign Profile** to assign a server profile to the selected PRIMERGY rack server or to the selected slot of a blade server. The **Select Profile** dialog box opens.

Unassign Profile

Click **Unassign Profile** to delete server profile assignments. The assignment is deleted and no server profile is assigned to the corresponding slot or PRIMERGY rack server. The display on the **Server Configuration** tab is updated accordingly.

Inventory Boot

Click **Inventory Boot** to re-create the inventory table of the selected server blade or PRIMERGY rack server.

During inventory boot the system BIOS assembles the inventory information of the server blade or PRIMERGY rack server hardware as needed by the Virtual-IO Manager and sends it to the management blade or to iRMC, where it is stored.

In some cases a manual execution of the inventory boot is necessary in order to support the new functionality:

- For support of new functionality a new version of the inventory table might be needed. Example: Converged network adapter card (CNA)
- In some cases the new version of the inventory table might not be created automatically for the server. Example: The new version of the inventory table is only supported by a new version of the system BIOS.

You must execute the inventory boot manually, after applying a new version of the system BIOS and iRMC firmware or a new version of the firmware of the optional hardware.

Boot

Click **Boot** to start the selected server blade(s) and PRIMERGY rack servers.

Shutdown

Click **Shutdown** to switch the selected server blades and PRIMERGY rack servers off. You can only assign a server profile to a server blade or PRIMERGY rack server if it is switched off (power off). In the next dialog box, you can select **Graceful Shutdown** or **Forced Power Off**.

Graceful Shutdown

Shuts down the server properly and then switches it off.

Forced Power Off
Switches off the server irrespective of the status of the operating system.

Update States
Click **Update States** to update the display in the **State** column for all server blades and PRIMERGY rack servers.

Video Redirection
Click **Video Redirection**. A new window opens where the console output of the server is shown.

Context menu
In the table or the graphic, several actions can be choosen from the context menu of the selected server (click right mouse button):

Assign Profile
Assign a server profile to a slot or PRIMERGY rack server.

Unassign Profile
Remove the assignment of a server profile from a slot or PRIMERGY rack server.

Create Profile
Create a server profile.

Show Profile Details
View definition of the assigned server profile.

Update State
Update the power state display, the boot mode and the virtualization status of a server.

Inventory Boot
Re-create the inventory table of the server blade(s) and PRIMERGY rack server(s).

Boot
Start the server blade(s) and PRIMERGY rack server(s).

Shutdown
Switch off the server blade(s) and PRIMERGY rack server(s).

Failover
Assign the server profile of a slot to a suitable spare slot in the event of a failure.

For more information on this, see chapter "Defining and assigning server profiles" on page 289".

5.3.5 Chassis Configuration tab

The **Chassis Configuration** tab provides an overview of the blade server configuration:

- The IBP configuration
- The server profile assignment

You can select an uplink port, a network or a bay by clicking it. The elements associated with the element you have selected are then highlighted.

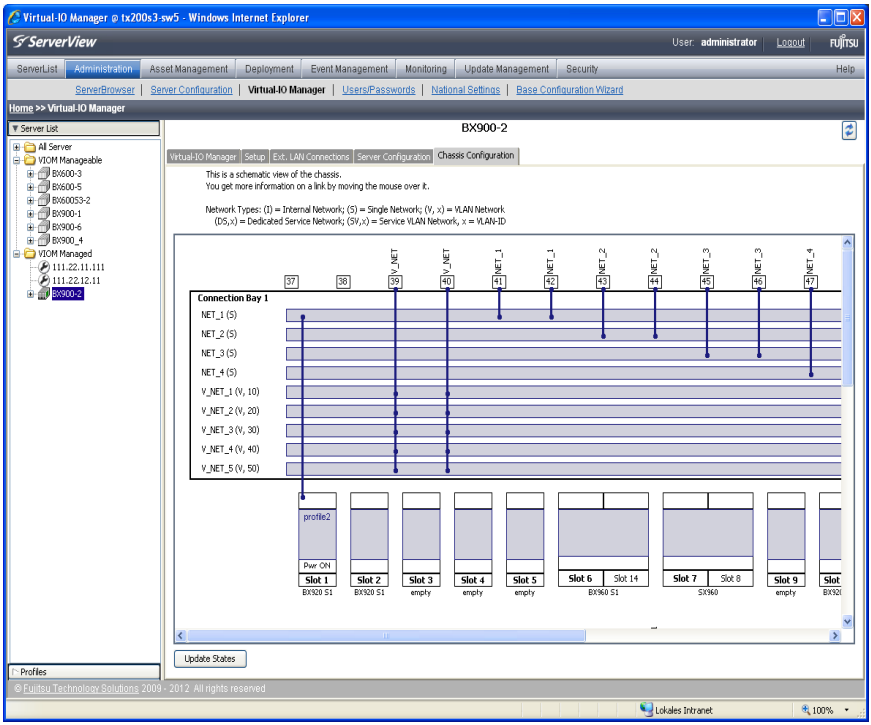

Figure 30: Chassis Configuration tab

The defined networks are listed in a rectangular box for each connection bay with an IBP connection blade. In the example above, connection bay 1 has the networks **NET_1**, **NET_2**, **NET_3**, **NET_4**, **V_NET_1**, …. Each network has a continuous gray row.

At the top edge of the rectangle, the uplink ports of the connection blade connected are displayed as squares and contain the port number. These are ports 37, 38, 39, …. The uplink set to which the port belongs is also listed. This means, for example, that ports 41 and 42 belong to the uplink set **NET_1**.

The blue lines from an uplink port to a network indicate which network is assigned to an uplink set. In the example above, network **NET_1** is therefore assigned to the uplink set **NET_1** with ports 41 and 42. The VLAN networks **V_NET_1**, …, **V_NET_5** belong to the uplink set **V_NET**.

The slots are displayed as smaller rectangles (**Slot 1**, **Slot 2**, **Slot 3**, …) below the rectangle that contains the list of networks. The name of an assigned server profile is entered in these rectangles. In the example above, **profile2** is entered for **Slot 1**.

Also in this rectangle, switched-on server blades are indicated with **Pwr ON**. You can update this display of the on/off status of the server blades with the **Update States** button.

The type of server blade connected (e.g., **BX920 S1** or **<empty>** if the slot is empty) is displayed below the slots.

The lines from a slot to a network indicate the network to which a server blade is connected. Tagged VLAN networks are marked as green lines with square endpoints, all other networks as blue lines with round endpoints. The slots are displayed for each connection bay.

In the figure above, only the connections for Connection Bay 1 are displayed. The other connection bays are displayed below this one. The server blade in **Slot 1** is connected to the network **NET_1** via the server profile **profile2**. This applies to LAN ports 1 and 3 of this blade (onboard LAN port 1 and 3 are assigned to Connection Bay 1).

Update States

Updates the display of the on/off status of the server blades in the rectangle of the slots.

5.3.6 Server Profiles view

The **Server Profiles** view shows the defined server profiles.

To open the **Server Profiles** view in the area on the right of the **ServerView Virtual-IO Manager** window, click **Profiles** in the area on the left of the **ServerView Virtual-IO Manager** window.

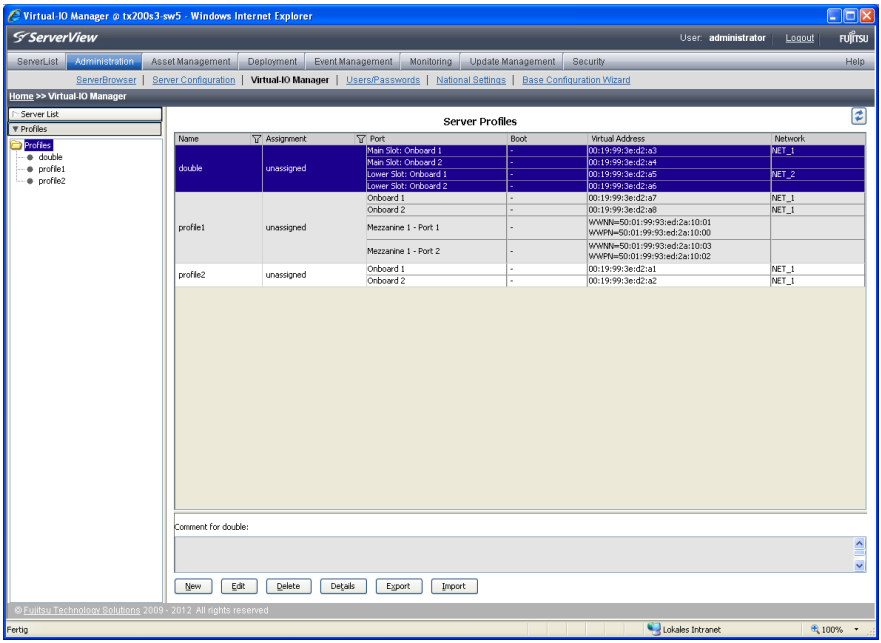

Figure 31: Server Profiles view

If in the **Server Profiles** view in the area on the left a folder is selected, all server profiles in this folder are shown in the table, otherwise the selected profile is shown. In the table there is one row per server profile showing the relevant information the profile:

Column	Significance
Name	The name os the server profile
Assignment	Chassis and slot or PRIMERGY rack server where this server profile is assigned, or **unassigned** if it is not assigned
Port	The identifier of a port; there is one row per port that is defined in the profile
Boot	Boot priority of this port, or "-" if it is no boot channel
Virtual Address	The virtual address of this port

Column	Significance
Network	The name of the network defining the dedicated LAN connection

Comment for
A more detailed description of the profile (optional).

Buttons

New
The **Create Server Profile** wizard is started to define a new profile.

Edit
The **Edit Server Profile** wizard is started to modify the selected profile. You can only modify server profiles that are not assigned. If more than one server profile is selected or the selected server profile is assigned, this button is disabled.

Delete
The selected profiles will be deleted. A message appears asking if you wish to delete the corresponding server profiles. If you confirm this, the server profiles will be deleted. You can only delete server profiles that are not assigned. If any selected server profile is assigned, this button is disabled.

Details
The **Serverprofile <profile name>** dialog box is opened providing information on the selected server profile.

Export
The selected server profiles are exported. A file selection box opens in which you select the name of the file to which you want to save the exported profiles.

Import
Server profiles are imported. In the file selection box that opens, select the file that you want to import.

Context menu
Several actions for a server profile can be choosen from the context menu (click right mouse button):

Edit Profile
The **Edit Server Profile** wizard is started to modify the selected profile. You can only modify server profiles that are not assigned. If more than one server profile is selected or the selected server profile is assigned, this menu item is disabled.

Show Profile Details
The **Serverprofile <profile name>** dialog box is opened providing information on the selected server profile. If more than one server profile is selected, this menu item is disabled.

Delete Profile
The selected profiles will be deleted. A message appears asking if you wish to delete the corresponding server profiles. If you confirm this, the server profiles will be deleted. You can only delete server profiles that are not assigned. If any selected server profile is assigned, this menu item is disabled.

5.4 Wizards

A wizard is an assistant that guides you, step by step, through a task.

A wizard usually consists of several steps that you work through in sequence. The number of steps used and their sequence are shown in a tree structure on the left. Steps that you have already completed are indicated in the tree.

The buttons in the bottom right of each step allow you to progress through the wizard workflow.

Previous
Opens the previous step in the wizard.

Next
Opens the next step in the wizard.

Finish
Executes the wizard with your settings.

Cancel
Cancels the wizard workflow without saving your changes.

Help
Launchs the context-sensitive online help.

Virtual-IO Manager provides the following wizards:

- **Create Network for IBP** wizard
- **Edit Uplink Set** wizard
- **Create Server Profile** wizard
- **Edit Server Profile** wizard
- **Configuration Backup/Restore** wizard

5.4.1 Create Network for IBP wizard (only for blade servers)

You use the **Create Network for IBP** wizard to define a network. In the Virtual-IO Manager, defining a network path refers to specifying which external ports are used to connect the relevant blade server chassis to which external networks.

Defining these types of network paths on an IBP module includes the following steps:

- Defining an uplink set. An uplink set contains a number of uplink ports. You can combine multiple uplink ports in one uplink set. You can configure the ports as active ports or as backup ports.
- Possibly also defining one or several meaningful network names, which are assigned to the uplink set.

The **Create Network for IBP** wizard comprises several dialog boxes to guide you through the individual steps. All required steps are displayed in the tree structure on the left.

You define network paths using the **Ext. LAN Connections** tab. This tab contains two other tabs (**Graphic** and **Details**). To open the **Create Network for IBP** wizard, click **New** on the **Graphic** or **Details** tab.

5.4.1.1 Select Type step (Create Network wizard)

Select Type is the first step in the **Create Network** wizard. In the first step, you specify what type of uplink set or what network you wish to create.

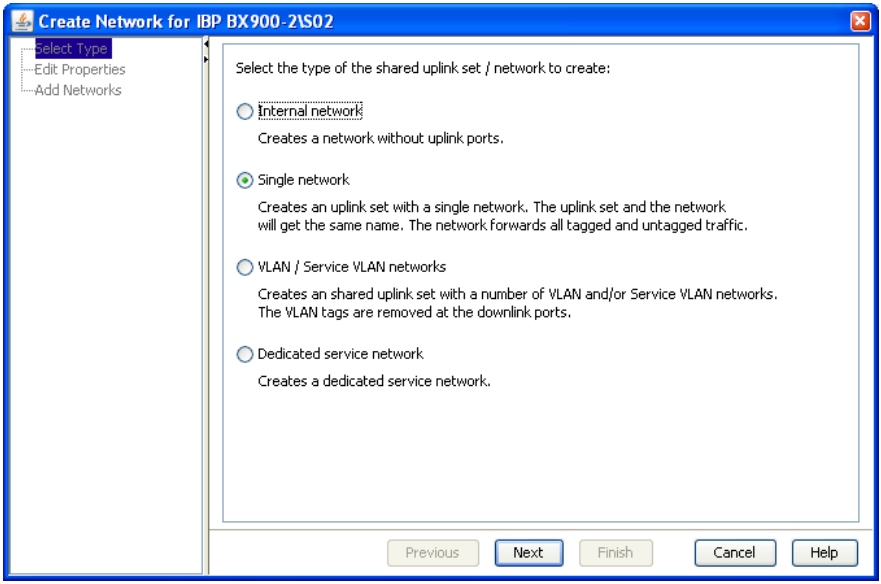

Figure 32: Create Network wizard (first step)

Internal network

Creates an internal network without a connection to an uplink port. This establishes a connection between the server blades (internal IBP connections) without there being a connection to an external network.

Single network (selected by default)
Creates an uplink set with a network. The uplink set and the network have the same name.

VLAN / Service VLAN networks
Creates an uplink set to which one network or even multiple networks with a VLAN ID can be assigned.

Dedicated Service network
Creates an uplink set with one dedicated service network.

5.4.1.2 Edit Properties step (Create Network wizard - internal network)

Edit Properties is the second step in the **Create Network** wizard.

The selection you made at the first step of the **Create Network** wizard determines which fields are displayed in the second step.

Edit Properties step for internal networks

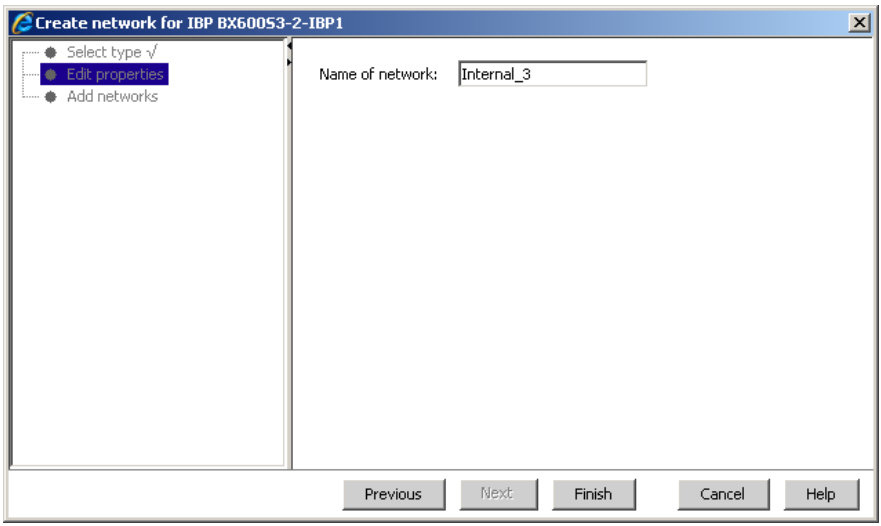

Figure 33: Create Network wizard (second step)

Name of network

Name for the internal network.

5.4.1.3 Edit Properties step (Create Network wizard - single /VLAN network)

Edit Properties is the second step in the **Create Network** wizard.

The selection you made at the first step of the **Create Network** wizard determines which fields are displayed in the second step.

Edit Properties step for single networks/VLAN networks

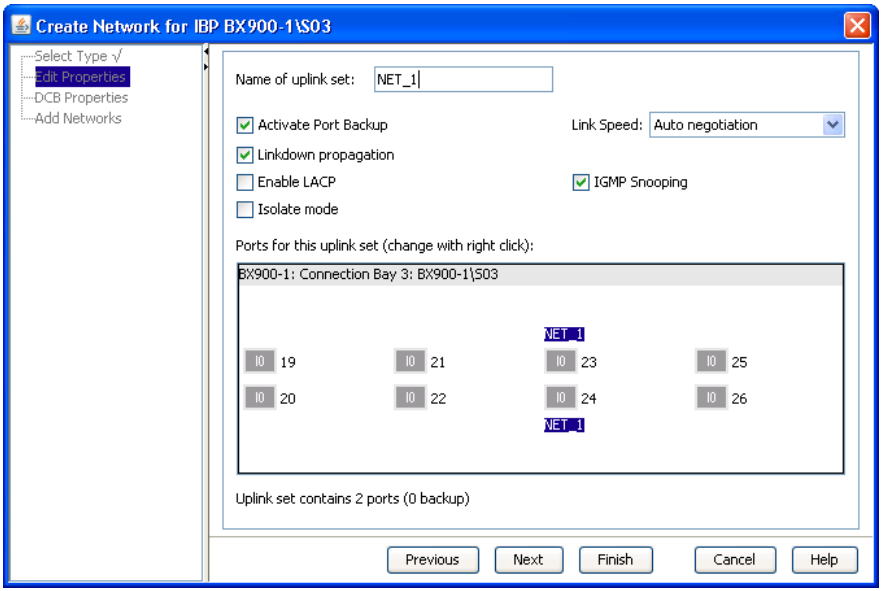

Figure 34: Create Network wizard (second step)

Name of uplink set

Name of uplink set. In the case of a single network, the network automatically has the same name.

Activate Port Backup

Switch to a backup port if an error occurs in the active port. The port backup function is only available if at least one backup port is

configured.

By default, **Activate Port Backup** is selected.

Linkdown propagation

Send a linkdown event if both the active ports and the backup ports fail. The linkdown event triggers a failover process on the server blade if configured accordingly.

If you select **Linkdown propagation**, the ports of the related server blades receive a linkdown event if a problem occurs. If the LAN drivers in the operating system of the server blade are configured accordingly, this linkdown event then triggers a failover on the second LAN port. A linkdown event is then triggered if all ports configured as active and all ports configured as backup ports of an uplink set fail. This allows the LAN connection to remain intact.

In order for the failover process in the server blade to work from LAN port 1 to LAN port 2, a LAN team must have been configured on the server blade and the network on the second IBP module must have been configured accordingly.

By default, **Linkdown propagation** is selected.

Enable LACP

You activate the LACP protocol via Enable LACP.

Isolate mode

You activate the isolate mode via **Isolate mode**. With isolate mode activated, server blades in the same chassis that use the same network cannot communicate with each other, but only with an external device though the uplinks.

This checkbox is not shown if the connection blade does not support isolate mode or if VLAN networks are created.

Link Speed

By selecting a **Link Speed**, you can select the transmission speed. If auto negotiation is not selected on the external switch, configure this value if a problem occurs using the drop-down menu according to the settings on the external switch. The following values are available:

Auto negotiation
The transmission speed is negotiated with the external switch. You can achieve a transmission speed of 1 Gbit/s with this value.

10 Mbit/s
10 Mbit/s full-duplex.

100 Mbit/s
100 Mbit/s full-duplex.

By default, **Auto negotiation** is selected.

IGMP Snooping
If **IGMP Snooping** is activated, the connection blade controls whether requests to join a multicast group occur at the downlink ports of an uplink set. If necessary, the corresponding downlink ports may be added to the forwarding table of this multicast group or removed again.

The **IGMP Snooping** option is selected by default.

Ports for the uplink set (change with right click)
Assign the required ports of the IBP to the uplink set. To do this, open the context menu of each relevant port. If an uplink set for a switch stack is defined, the uplink ports of all IBPs belonging to the stack are shown one below the other. Use the scroll bar to reach the uplink ports of another stack member.

The context menu can contain the following menu items:

Add
Only activated if the port is not yet assigned to the uplink set.

Assigns the port to the uplink set. If the port is already assigned to another uplink set, this assignment is deleted.

Add as Backup
Only activated if the port is not yet assigned to the uplink set.

Assigns the port as a backup port to the uplink set. If the port is already assigned to another uplink set, this assignment is deleted.

Remove

Only available if the port is already assigned to the uplink.

Removes the port assignment from the uplink set.

Active

Only available if the port is configured as **Backup**.

The corresponding port is configured as active and no longer as backup.

Backup

Only available if the port is configured as **Active**.

The corresponding port is configured as a backup port and no longer as active port.

Ports configured as active or backup ports are indicated by different colors in the display (light red/light green or dark red/dark green). The name of the uplink set is assigned to configured ports, which you can recognize by this.

With PY CB Eth Switch/IBP 1Gb 36/8+2 connection blades, 1GB uplinks and 10 Gb uplinks must not be mixed. If you try this, the **Next** and **Finish** buttons will not be enabled.

The **Next** and **Finish** buttons are disabled while not at least one active port is included in the uplink set.

5.4.1.4 Edit Properties step (Create Network wizard - dedicated service network)

Edit Properties is the second step in the **Create Network** wizard.

The selection you made at the first step of the **Create Network** wizard determines which fields are displayed in the second step.

Edit Properties step for dedicated service networks

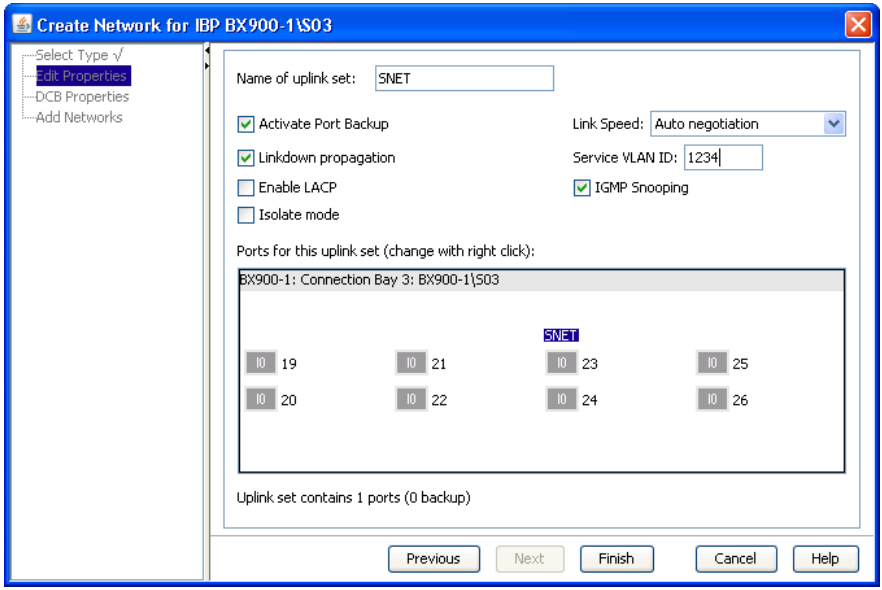

Figure 35: Create Network wizard (second step)

Name of uplink set

Name of uplink set. In the case of a single network, the network automatically has the same name.

Activate Port Backup

Switch to a backup port if an error occurs in the active port. The port backup function is only available if at least one backup port is configured.

By default, **Activate Port Backup** is selected.

Linkdown propagation

Send a linkdown event if both the active ports and the backup ports fail. The linkdown event triggers a failover process on the server blade if configured accordingly.

If you select **Linkdown propagation**, the ports of the related server blades receive a linkdown event if a problem occurs. If the LAN drivers in the operating system of the server blade are configured accordingly, this linkdown event then triggers a failover on the second LAN port. A linkdown event is then triggered if all ports configured as active and all ports configured as backup ports of an uplink set fail. This allows the LAN connection to remain intact.

In order for the failover process in the server blade to work from LAN port 1 to LAN port 2, a LAN team must have been configured on the server blade and the network on the second IBP module must have been configured accordingly.

By default, **Linkdown propagation** is selected.

Enable LACP

You activate the LACP protocol via Enable LACP.

Isolate mode

You activate the isolate mode via **Isolate mode**. With isolate mode activated, server blades in the same chassis that use the same network cannot communicate with each other, but only with an external device though the uplinks.

This checkbox is not shown if the connection blade does not support isolate mode.

Link Speed

By selecting a **Link Speed**, you can select the transmission speed. If auto negotiation is not selected on the external switch, configure this value if a problem occurs using the drop-down menu according to the settings on the external switch. The following values are available:

Auto negotiation

The transmission speed is negotiated with the external switch. You can achieve a transmission speed of 1 Gbit/s with this value.

10 Mbit/s

10 Mbit/s full-duplex.

100 Mbit/s
100 Mbit/s full-duplex.

By default, **Auto negotiation** is selected.

Service VLAN id
A Service VLAN ID must be specified.

IGMP Snooping
If **IGMP Snooping** is activated, the connection blade controls whether requests to join a multicast group occur at the downlink ports of an uplink set. If necessary, the corresponding downlink ports may be added to the forwarding table of this multicast group or removed again.

The **IGMP Snooping** option is selected by default.

Ports for the uplink set (change with right click)
Assign the required ports of the IBP to the uplink set. To do this, open the context menu of each relevant port. If an uplink set for a switch stack is defined, the uplink ports of all IBPs belonging to the stack are shown one below the other. Use the scroll bar to reach the uplink ports of another stack member.

The context menu can contain the following menu items:

Add
Only activated if the port is not yet assigned to the uplink set.

Assigns the port to the uplink set. If the port is already assigned to another uplink set, this assignment is deleted.

Add as Backup
Only activated if the port is not yet assigned to the uplink set.

Assigns the port as a backup port to the uplink set. If the port is already assigned to another uplink set, this assignment is deleted.

Remove
Only available if the port is already assigned to the uplink.

Removes the port assignment from the uplink set.

Active

Only available if the port is configured as **Backup**.

The corresponding port is configured as active and no longer as backup.

Backup

Only available if the port is configured as **Active**.

The corresponding port is configured as a backup port and no longer as active port.

Ports configured as active or backup ports are indicated by different colors in the display (light red/light green or dark red/dark green). The name of the uplink set is assigned to configured ports, which you can recognize by this.

With PY CB Eth Switch/IBP 1Gb 36/8+2 connection blades, 1GB uplinks and 10 Gb uplinks must not be mixed. If you try this, the **Next** and **Finish** buttons will not be enabled.

The **Next** and **Finish** buttons are disabled while not at least one active port is included in the uplink set.

5.4.1.5 DCB Properties step (Create Network wizard - single/VLAN network)

DCB Properties is the third step in the **Create Network** wizard if a single network or a VLAN network is created.

This step is only displayed if the connection blade supports DCB (Data Center Bridging), i.e. for PY CB Eth Switch/IBP 10 Gb 18/8 connection blades.

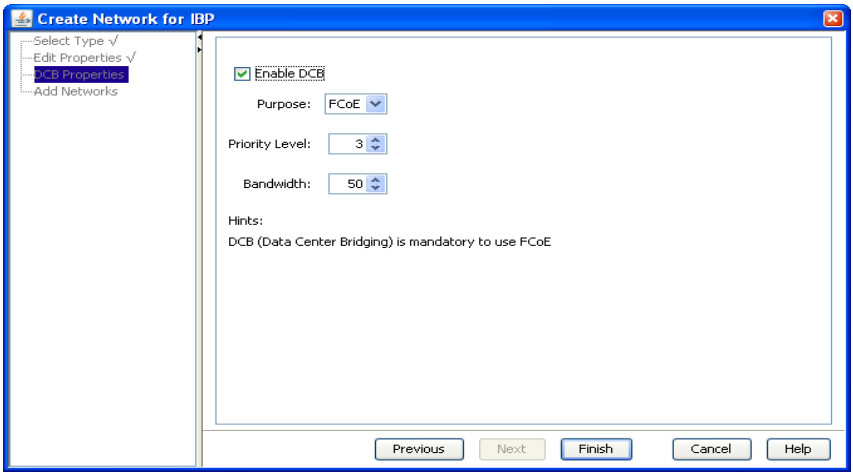

Figure 36: Create Network wizard (DCB Properties)

Enable DCB

Enables the DCB feature of the connection blade. The DCB (Data Center Bridging) settings here are specific configuration settings in a DCB-enabled switching device.

This option should be enabled if this uplink set should be used for FCoE (Fibre Channel over Ethernet).

Purpose

The purpose of the DCB setting. Possible values are FCoE and iSCSI.

Priority Level

The priority level. Possible values are 0 to 7; for FCoE this is by default the value 3, for iSCSI the default value is 4.

Bandwidth

The share of the bandwidth in percent that is assigned to this function. If the sum of all bandwidths of one IO-channel is not 100, the values are internally adjusted accordingly.

This is the bandwidth reserved for the FCoE function. The FCoE function might share the complete bandwidth of 10 Gb with other functions. A value of 60, for example, means that a bandwidth of at least 6 Gb/sec is reserved for the FCoE packages.

5.4.1.6 Add Networks step (Create Network wizard - VLAN network)

Add Networks is the third step in the **Create Network** wizard if a VLAN network is created.

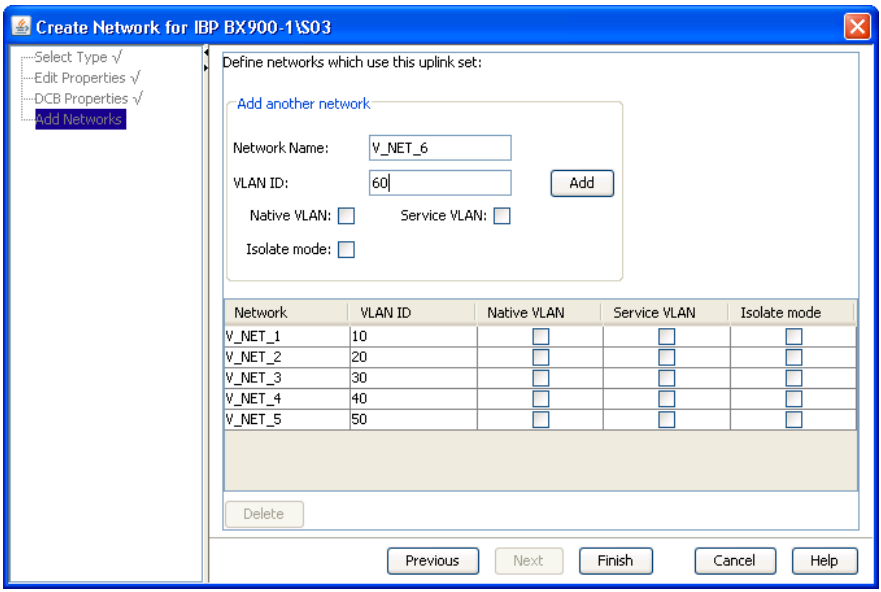

Figure 37: Create Network wizard (third step)

In this step, you assign meaningful names for the networks defined via VLAN IDs, which are to be assigned to the uplink set defined in the previous step.

Add another network

Network Name

meaningful name of the network.

VLAN-Id
The unique VLAN number of the network specified above. The VLAN ID must be assigned uniquely within a shared uplink set.

Native VLAN
Defines a network as native VLAN. All packages that do not contain a VLAN ID will be allowed through this connection.

Service VLAN
Defines a network as service VLAN.

Isolate mode
You activate the isolate mode via **Isolate mode**. With isolate mode activated, server blades in the same chassis that use the same network cannot communicate with each other, but only with an external device though the uplinks.

This checkbox is not shown if the connection blade does not support isolate mode.

Add button
Adds the virtual LAN network to the table below.

Table
You can also define a network as a native VLAN or service VLAN and set the isolate mode in the table retrospectively:

Native VLAN
Select the checkbox in the corresponding row under **Native VLAN** in the table to define a network as a native VLAN.

Service VLAN
Select the checkbox in the corresponding row under **Service VLAN** in the table to define a network as a service VLAN.

Isolate mode
Select the checkbox in the corresponding row under **Isolate mode** to set the isolate mode for this network. With isolate mode activated, servers

in the same chassis that use the same network cannot communicate with each other, but only with an external device though the uplinks.

The column **Isolate mode** is not shown if the connection blade does not support isolate mode.

Delete

Deletes a selected VLAN networks from the list.

5.4.2 Edit Uplink Set wizard

You use the **Edit Uplink Set** wizard to modify an defined uplink set.

The **Edit Uplink Set** wizard comprises several dialog boxes to guide you through the individual steps. All required steps are displayed in the tree structure on the left.

You modify an defined uplink using the **Ext. LAN Connections** tab. This tab contains two other tabs (**Graphic** and **Details**). To open the **Edit Uplink Set** wizard, click **Edit** on the **Graphic** or **Details** tab.

5.4.2.1 Edit Properties step (Edit Uplink Set wizard - single/VLAN network)

Edit Properties is the first step in the **Edit Uplink Set** wizard.

The type of the network determines which fields are displayed in this step.

Edit Properties step for single networks/VLAN networks

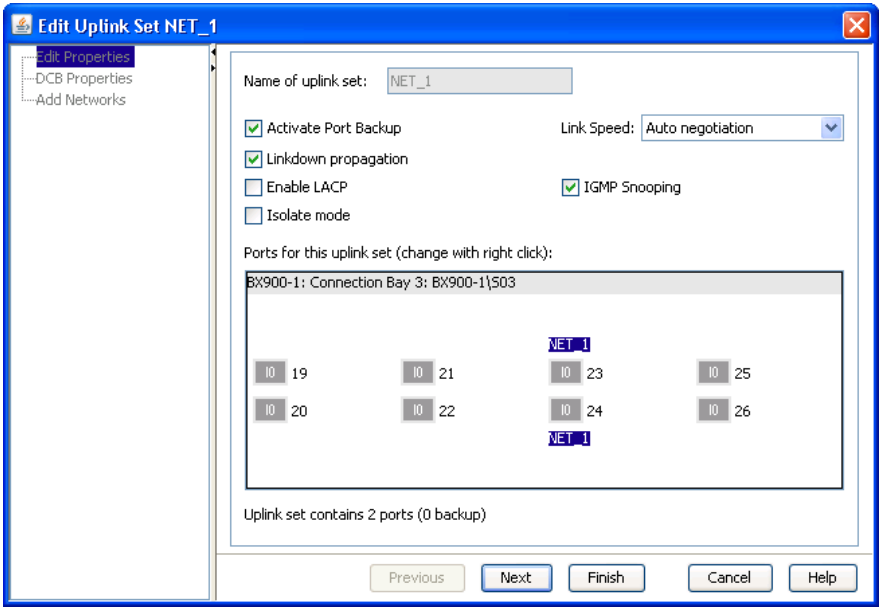

Figure 38: Edit Uplink Set wizard (first step)

You can modify following parameters in the first step of the wizard:

Activate Port Backup

Switch to a backup port if an error occurs in the active port. The port backup function is only available if at least one backup port is configured.

Linkdown propagation

Send a linkdown event if both the active ports and the backup ports fail. The linkdown event triggers a failover process on the server blade if configured accordingly.

If you select **Linkdown propagation**, the ports of the related server blades receive a linkdown event if a problem occurs. If the LAN drivers in the operating system of the server blade are configured accordingly, this linkdown event then triggers a failover on the second LAN port. A

linkdown event is then triggered if all ports configured as active and all ports configured as backup ports of an uplink set fail. This allows the LAN connection to remain intact.

In order for the failover process in the server blade to work from LAN port 1 to LAN port 2, a LAN team must have been configured on the server blade and the network on the second IBP module must have been configured accordingly.

Enable LACP

You activate the LACP protocol via Enable LACP.

Isolate mode

You activate the isolate mode via **Isolate mode**. With isolate mode activated, server blades in the same chassis that use the same network cannot communicate with each other, but only with an external device though the uplinks.

This checkbox is not shown if the connection blade does not support isolate mode or if VLAN networks are edited.

Link Speed

By selecting a **Link Speed**, you can select the transmission speed. If auto negotiation is not selected on the external switch, configure this value if a problem occurs using the drop-down menu according to the settings on the external switch. The following values are available:

Auto negotiation

The transmission speed is negotiated with the external switch. You can achieve a transmission speed of 1 Gbit/s with this value.

10 Mbit/s

10 Mbit/s full-duplex.

100 Mbit/s

100 Mbit/s full-duplex.

IGMP Snooping

If **IGMP Snooping** is activated, the connection blade controls whether requests to join a multicast group occur at the downlink ports of an uplink set. If necessary, the corresponding downlink ports may be added to the forwarding table of this multicast group or removed again.

Ports for the uplink set (change with right click)

Assign the required ports of the IBP to the uplink set. To do this, open the context menu of each relevant port. If an uplink set for a switch stack is defined, the uplink ports of all IBPs belonging to the stack are shown one below the other. Use the scroll bar to reach the uplink ports of another stack member.

The context menu can contain the following menu items:

Add

Only activated if the port is not yet assigned to the uplink set.

Assigns the port to the uplink set. If the port is already assigned to another uplink set, this assignment is deleted.

Add as Backup

Only activated if the port is not yet assigned to the uplink set.

Assigns the port as a backup port to the uplink set. If the port is already assigned to another uplink set, this assignment is deleted.

Remove

Only available if the port is already assigned to the uplink.

Removes the port assignment from the uplink set.

Active

Only available if the port is configured as **Backup**.

The corresponding port is configured as active and no longer as backup.

Backup

Only available if the port is configured as **Active**.

The corresponding port is configured as a backup port and no longer as active port.

Ports configured as active or backup ports are indicated by different colors in the display (light red/light green or dark red/dark green). The name of the uplink set is assigned to configured ports, which you can recognize by this.

5.4.2.2 Edit Properties step (Edit Uplink Set wizard - dedicated service network)

Edit Properties is the first step in the **Edit Uplink Set** wizard.

The type of the network determines which fields are displayed in this step.

Edit Properties step for dedicated service networks

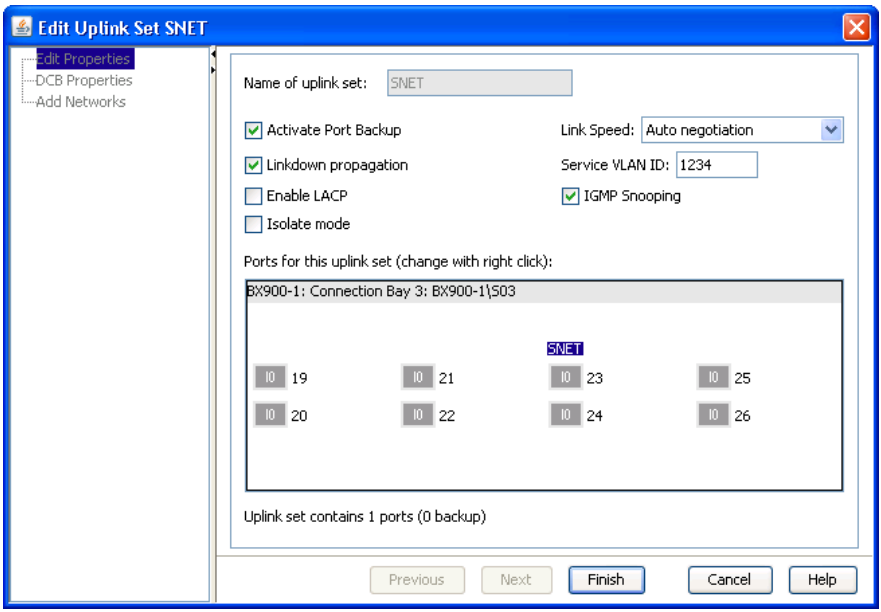

Figure 39: Edit Uplink Set wizard (first step)

You can modify following parameters in the first step of the wizard:

Activate Port Backup

Switch to a backup port if an error occurs in the active port. The port backup function is only available if at least one backup port is configured.

Linkdown propagation

Send a linkdown event if both the active ports and the backup ports fail. The linkdown event triggers a failover process on the server blade if configured accordingly.

If you select **Linkdown propagation**, the ports of the related server blades receive a linkdown event if a problem occurs. If the LAN drivers in the operating system of the server blade are configured accordingly, this linkdown event then triggers a failover on the second LAN port. A linkdown event is then triggered if all ports configured as active and all ports configured as backup ports of an uplink set fail. This allows the LAN connection to remain intact.

In order for the failover process in the server blade to work from LAN port 1 to LAN port 2, a LAN team must have been configured on the server blade and the network on the second IBP module must have been configured accordingly.

Enable LACP

You activate the LACP protocol via Enable LACP.

Isolate mode

You activate the isolate mode via **Isolate mode**. With isolate mode activated, server blades in the same chassis that use the same network cannot communicate with each other, but only with an external device though the uplinks.

This checkbox is not shown if the connection blade does not support isolate mode.

Link Speed

By selecting a **Link Speed**, you can select the transmission speed. If auto negotiation is not selected on the external switch, configure this

value if a problem occurs using the drop-down menu according to the settings on the external switch. The following values are available:

Auto negotiation
The transmission speed is negotiated with the external switch. You can achieve a transmission speed of 1 Gbit/s with this value.

10 Mbit/s
10 Mbit/s full-duplex.

100 Mbit/s
100 Mbit/s full-duplex.

Service VLAN id
A Service VLAN ID must be specified.

IGMP Snooping
If **IGMP Snooping** is activated, the connection blade controls whether requests to join a multicast group occur at the downlink ports of an uplink set. If necessary, the corresponding downlink ports may be added to the forwarding table of this multicast group or removed again.

Ports for the uplink set (change with right click)
Assign the required ports of the IBP to the uplink set. To do this, open the context menu of each relevant port. If an uplink set for a switch stack is defined, the uplink ports of all IBPs belonging to the stack are shown one below the other. Use the scroll bar to reach the uplink ports of another stack member.

The context menu can contain the following menu items:

Add
Only activated if the port is not yet assigned to the uplink set.

Assigns the port to the uplink set. If the port is already assigned to another uplink set, this assignment is deleted.

Add as Backup
Only activated if the port is not yet assigned to the uplink set.

Assigns the port as a backup port to the uplink set. If the port is already assigned to another uplink set, this assignment is deleted.

Remove

Only available if the port is already assigned to the uplink.

Removes the port assignment from the uplink set.

Active

Only available if the port is configured as **Backup**.

The corresponding port is configured as active and no longer as backup.

Backup

Only available if the port is configured as **Active**.

The corresponding port is configured as a backup port and no longer as active port.

Ports configured as active or backup ports are indicated by different colors in the display (light red/light green or dark red/dark green). The name of the uplink set is assigned to configured ports, which you can recognize by this.

5.4.2.3 DCB Properties (Edit Uplink Set wizard - single/VLAN network)

DCB Properties is the second step in the **Edit Uplink Set** wizard if a single network or a VLAN network is edited.

It is only displayed if the connection blade supports DCB (Data Center Bridging), i.e. for PY CB Eth Switch/IBP 10 Gb 18/8 connection blades.

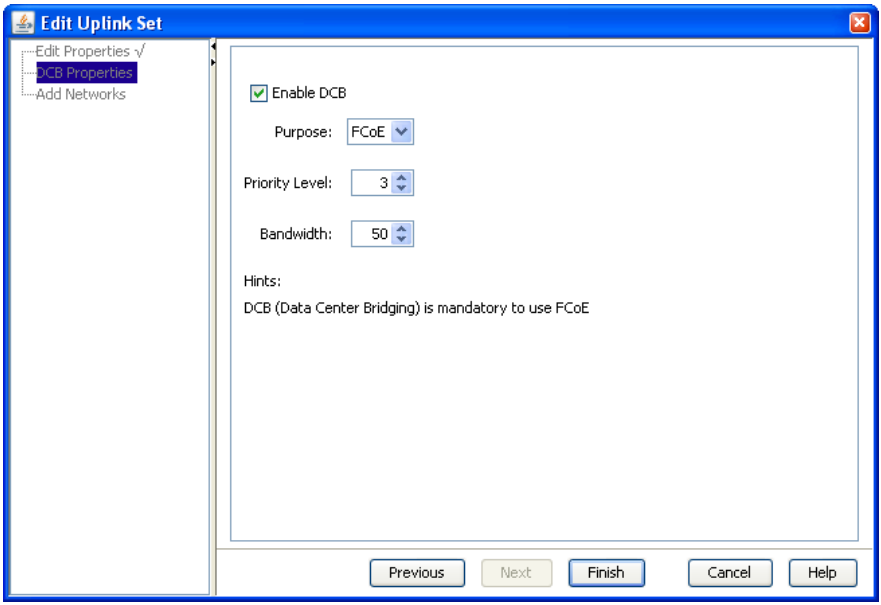

Figure 40: DCB Properties

Enable DCB

Enables the DCB feature of the connection blade. The DCB (Data Center Bridging) settings here are specific configuration settings in a DCB-enabled switching device.

This option should be enabled if this uplink set should be used for FCoE (Fibre Channel over Ethernet).

Purpose

The purpose of the DCB setting. Possible values are FCoE and iSCSI.

Priority Level

The priority level. Possible values are 0 to 7; for FCoE this is by default the value 3, for iSCSI the default value is 4.

Bandwidth

The share of the bandwidth in percent that is assigned to this function. If the sum of all bandwidths of one IO-channel is not 100, the values are

internally adjusted accordingly.

This is the bandwidth reserved for the FCoE function. The FCoE function might share the complete bandwidth of 10 Gb with other functions. A value of 60, for example, means that a bandwidth of at least 6 Gb/sec is reserved for the FCoE packages.

5.4.2.4 Add Networks step (Edit Uplink Set wizard - VLAN network)

Add Networks is the second step in the **Edit Uplink Set** wizard if a VLAN network is edited.

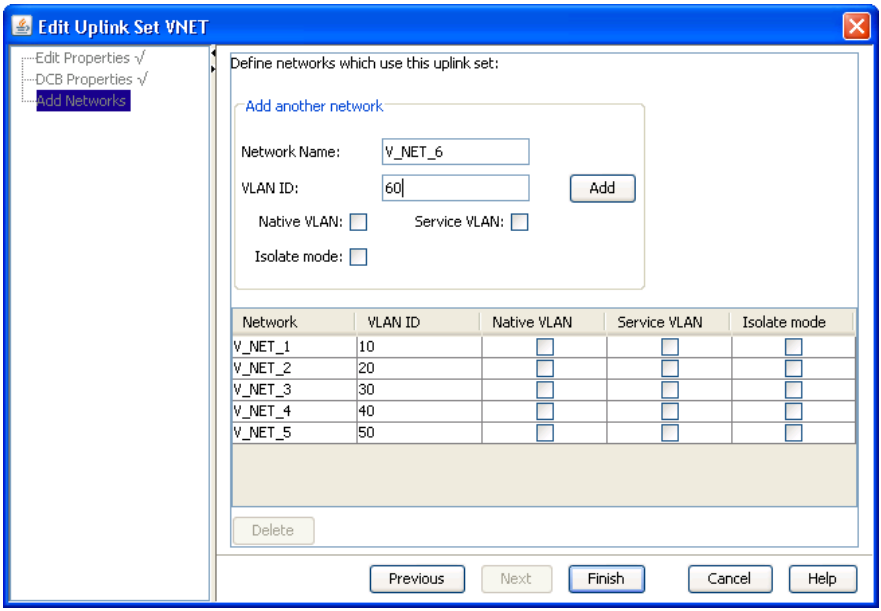

Figure 41: Edit Uplink Set wizard (second step)

In this step, you assign meaningful names for the networks defined via VLAN IDs, which are to be assigned to the uplink set defined in the previous step.

Add another network

Network Name
meaningful name of the network.

VLAN Id
The unique VLAN number of the network specified above. The VLAN ID must be assigned uniquely within a shared uplink set.

Native VLAN
Defines a network as native VLAN. All packages that do not contain a VLAN ID will be allowed through this connection.

Service VLAN
Defines a network as service VLAN.

Isolate mode
You activate the isolate mode via **Isolate mode**. With isolate mode activated, server blades in the same chassis that use the same network cannot communicate with each other, but only with an external device though the uplinks.

This checkbox is not shown if the connection blade does not support isolate mode.

Add button
Adds the virtual LAN network to the table below.

Table
You can also define a network as a native VLAN or service VLAN and set the isolate mode in the table retrospectively:

Native VLAN
Select the checkbox in the corresponding row under **Native VLAN** in the table to define a network as a native VLAN.

Service VLAN
Select the checkbox in the corresponding row under **Service VLAN** in the table to define a network as a service VLAN.

Isolate mode
Select the checkbox in the corresponding row under **Isolate mode** to set the isolate mode for this network. With isolate mode activated, servers in the same chassis that use the same network cannot communicate with each other, but only with an external device though the uplinks.

The column **Isolate mode** is not shown if the connection blade does not support isolate mode.

Delete button
Deletes a selected VLAN networks from the list.

5.4.3 Create Server Profile wizard

You use this wizard to define a new server profile.

The **Create Server Profile** wizard comprises several dialog boxes to guide you through the individual steps. All required steps are displayed in the tree structure on the left.

You define a server profile in the **Server Profile** view in the area on the right of the **ServerView Virtual-IO Manager** window. To open the **Create Server Profile** wizard, click the **New** button in the area on the right or select **New Profile** from the **Profiles** group context menu.

Another way to define server profiles is to switch to the view of the servers managed by VIOM in the area on the left using the **Server List** button. Then switch to the **Server Configuration** tab on the right and select **Create Profile** from the context menu of a table entry.

5.4.3.1 Name step (Create Server Profile wizard)

Name is the first step in the **Create Server Profile** wizard. In the first step, you specify the name of the server profile, its intended target, and optionally a comment.

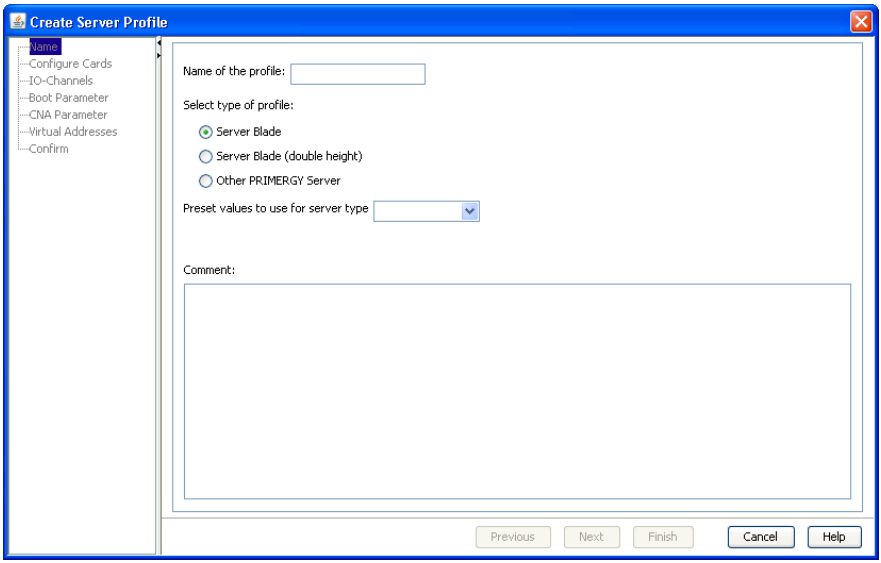

Figure 42: Name step

Name of the profile

Name of the server profile. If a profile already exists with this name or the name is invalid, the name is marked in red.

Select type of profile

The target type to which this profile should be assignable:

Server Blade

This profile can be assigned to slots of a blade server.

Server Blade (double height)

This profile can be assigned to two slots of a BX900 which are one over the other.

Other PRIMERGY Server

The profile can be assigned to a PRIMERGY rack server.

Preset values to use for server type

Server model (optional).

The number and type of LAN ports under **Onboard IO channels** are adjusted automatically according to the server model you select. The number of LAN ports and mezzanine/PCI cards cannot exceed the maximum possible value for the selected server model.

Comment

Comments on a more detailed description of the profile (optional)

5.4.3.2 Configure Cards step (Create Server Profile wizard)

Configure Cards is the next step in the **Create Server Profile** wizard.

In this step, you specify the number of mezzanine/PCI cards as well as the number and type of I/O channels for each card and for onboard.

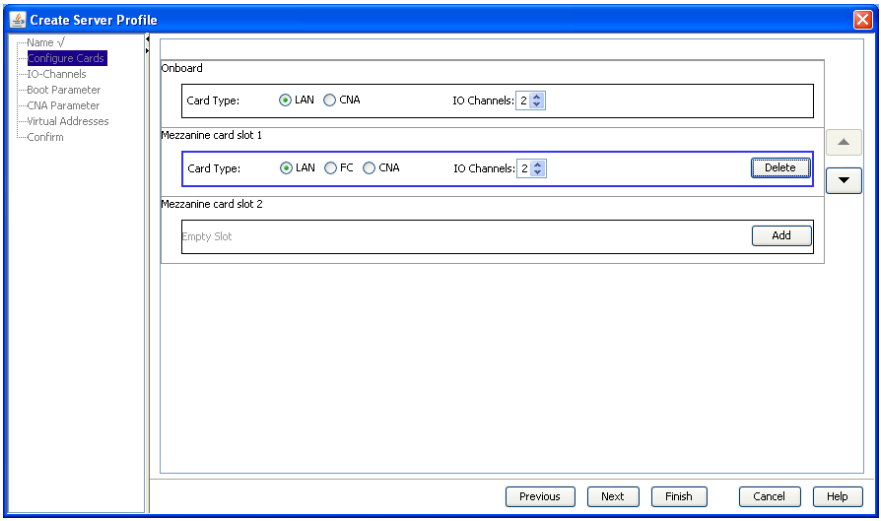

Figure 43: Configure Card step

In the **Onboard** box you can specify the type and number of onboard ports to be supported with this server profile. If you selected a server blade model under **Preset values to use for server type** in the **Name** step, the number and type of ports is initialized according to the model. You can adjust the number for the server profile. However, you cannot exceed the maximum possible value for the selected server model.

Depending on the selected profile and server type, there are a number of boxes, each representing a slot for a mezzanine or PCI card. If you want to configure a mezzanine or PCI card with this server profile, you must click the **Add** button in the corresponding slot to add a card. For each card you must specify its type:

LAN
LAN card. This card can have up to four ports.

FC
Fibre Channel card. This card can have up to two ports.

CNA
CNA card. This card can have up to two ports.

With **IO Channels** you specify the number of ports for the card to be supported with this server profile.

If you do not use an already configured mezzanine card, you can remove it with the **Delete** button.

With the arrow buttons on the right you can change the position of a mezzanine/PCI card. To do this, you must select a mezzanine/PCI card by clicking in its box (a blue border indicates that it is selected). You can use the up and down arrows to move the selected card by one position. If there is already a card in the target position, the cards swap their places. Any configuration for the mezzanine/PCI cards (e.g. boot configuration or virtual addresses) is retained.

5.4.3.3 IO-Channels step (Create Server Profile wizard)

IO-Channels is the next step in the **Create Server Profile** wizard.

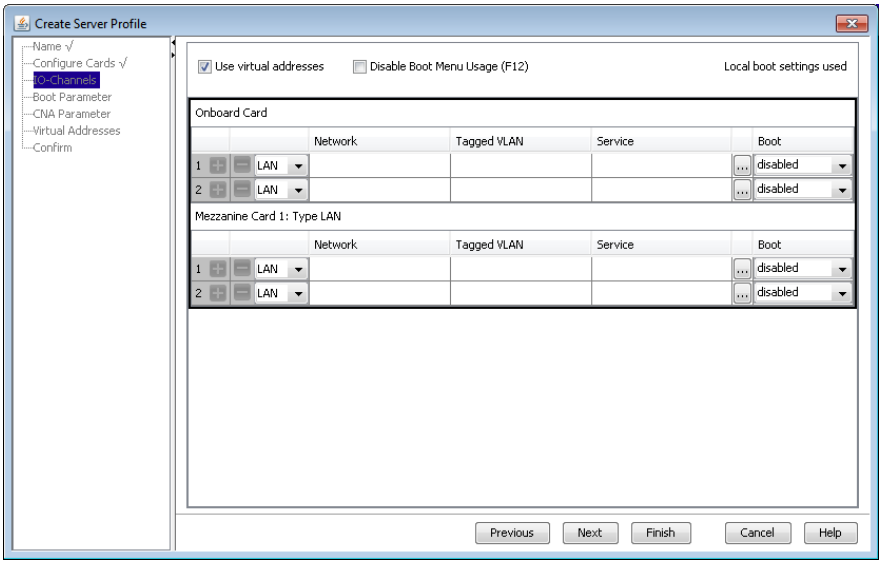

Figure 44: IO-Channels step

Local boot settings used

This message shows that the local boot settings are being used in the server profile. This message is not displayed if a boot device is used in the server profile.

Use virtual addresses

Uses virtual MAC addresses and WWN addresses with this profile. You can enter this information in another step within the wizard or VIOM can assign this information automatically.

Disable Boot Menu Usage (F12)

Prevents the VIOM boot settings from being overwritten on your local computer.

SMUX setting

This is only visible if a second mezzanine card of a blade server profile is defined as a LAN mezzanine card.

Defines the fabric to which the second mezzanine card is routed:

Fabric 3
All paths are routed to fabric 3.

Fabric 4
All paths are routed to fabric 4.

Fabric 3 & 4
LAN1 is routed to fabric 3 and LAN2 to fabric 4.

For more information on SMUX settings, see the documentation "PRIMERGY BX900 Blade Server Systems - ServerView Management Blade S1".

The upper table displays the onboard ports (up to 6). Another table is displayed for each available mezzanine/PCI card configured in the previous step. The tables have the following columns:

Column	Significance
first column	Port number of the onboard or the mezzanine/PCI card port. On CNA cards the type of the function (LAN, FCoE, iSCSI) can be selected from the drop-down list. Another function can be added to this port with the  button. Functions can be removed with the  button.

Column	Significance
Network	Name of the network. You can specify a network for each LAN or CNA port, but not for profiles of type **Rack Server**. If you want to use the profile on blade servers with IBP modules, you can specify a network for each LAN or CNA port. If you work with blade servers that have non-VIOM-capable LAN modules (Open Fabric mode), do not specify a network as it is not possible to define networks on these modules. To enter a network name, click the table cell to switch to edit mode. You can also open a network selection dialog box via the "..." button. In it you can select a managed blade server chassis from the selection list. The networks defined for this chassis that are suitable for the selected port are then displayed. To select a network, double-click the name or select the name and click the **Add** button. Make sure that the networks entered here are/will be configured before the profile is activated on the corresponding blade server. As long as a network does not yet exist, the server profile can be created with this network, but cannot yet be assigned to a slot. If you wish to exit the network selection without selecting a network, click the **Close** button or another input field.
Tagged VLAN	Names of tagged VLAN networks. You can specify tagged VLAN networks for each LAN or CNA port, but not for profiles of type **Rack Server**. If you specify more than one tagged VLAN for a port, the names must be separated by commas. If you use the network selection box, the name of the chosen network is added to the **Tagged VLAN** column if you use the **Add tagged** button.

Column	Significance
Service	Names of the service networks. You can specify service networks for each LAN or CNA port, but not for profiles of type **Rack Server**. If you specify more than one service network for a port, the names must be separated by commas. If you use the network selection box, the name of the chosen network is replaced in the **Network** column or added to the **Service** column depending on the type of the selected network.
Boot	To configure the port as the boot device, select **PXE boot**, **iSCSI boot**, or **SAN boot** from the selection list. If you select **disabled**, this port is not a boot device. You can define up to four boot devices in a profile. The values available in the selection list depends on type of port or function. If you configure an iSCSI boot device or SAN boot device, you must specify additional boot settings in the **Boot Parameter** step.

5.4.3.4 Boot Parameter step (Create Server Profile wizard)

Boot Parameter is the next step in the **Create Server Profile** wizard. In this step, you specify the boot device and boot parameter order. This step is only shown if you configured a iSCSI boot or SAN boot or at least two channels as boot devices in the **IO-Channels** step.

If you configured at least two channels as boot devices in the previous step, use the arrow up and arrow down buttons to change the boot order.

The selection you made at the **IO-Channels** step (iSCSI boot or SAN boot) determines which fields are displayed in this dialog box.

Boot Parameters for an iSCSI boot (LAN ports or iSCSI function on CNA mezzanine card)

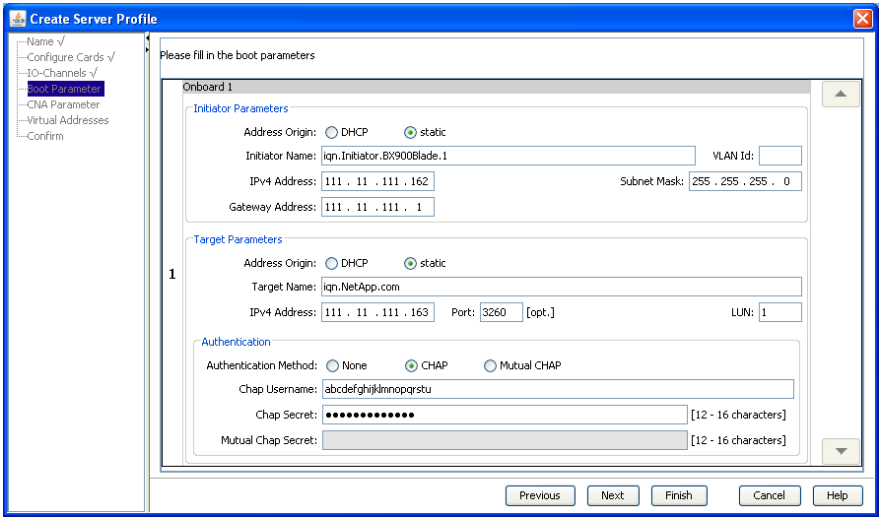

Figure 45: Boot Parameter step (iSCSI boot)

Initiator Parameters

Address Origin

DHCP

The system tries (in the case of an iSCSI boot) to obtain the client IP address, subnet mask, and gateway IP address from a DHCP server. Only the initiator name and (optional) the VLAN ID must be specified here.

static

A static client IP address, subnet mask, and gateway IP address must be specified.

Initiator Name

Name of the iSCSI initiator to be used (in the case of an iSCSI boot) for the connection to the iSCSI target.

VLAN Id

VLAN ID that is used by the HBA to send its requests.(optional)

Should be used only for CNA-iSCSI-functions.

IPv4 address

Static client IP address to be used for this port. The port will use this IP address for the entire iSCSI session. You can enter an IP address in this field if **Address Origin: static** is selected.

Subnet Mask

IP subnet mask. This should be the IP subnet mask of the network used to connect this port (in the case of an iSCSI boot). You can enter the subnet mask in this field if **Address Origin: static** is selected.

Gateway Address

IP address of the network gateway. You can enter the gateway address in this field if **Address Origin: static** is selected. This is necessary if the iSCSI target is in a subnetwork other than the subnetwork of the selected iSCSI boot port.

If iSCSI initiator and target are in the same network segment, no gateway is needed. The gateway address should be set to the value **0.0.0.0**.

Target Parameters

Address Origin

DHCP

The system tries (in the case of an iSCSI boot) to obtain the name of the iSCSI target, the IP address of the iSCSI target, the IP port number, and the SCSI LUN ID from a DHCP server in the network.

static

Static name for the iSCSI target, a static IP address for the iSCSI target, a static IP port number, and a static SCSI LUN ID.

Target Name
IQN name of the iSCSI target. You can enter a name in this field if **Address Origin: static** is selected.

IPv4 address
IP address of the iSCSI target. You can enter an IP address in this field if **Address Origin: static** is selected.

Port
TCP port number (default: 3260 for iSCSI). You can enter a port number in this field if **Address Origin: static** is selected. (optional)

LUN
LUN ID of the boot disk on the SCSI target. You can enter a LUN ID in this field if **Address Origin: static** is selected.

Authentication

Authentication Method

None
No authentication is used.

CHAP
CHAP authentication is activated for this port. CHAP allows the target to authenticate the initiator. After activating CHAP, you must enter a user name and password for the target.

Mutual CHAP
Mutual CHAP authentication is activated for this port. Mutual CHAP allows the initiator to authenticate the target. After activating mutual CHAP authentication, you must enter a user name, a password for the target, and a mutual CHAP password.

Chap Username
CHAP user name. The name must be identical to the name configured on the iSCSI target.

Chap Secret

CHAP password. This password must be identical to the password configured on the iSCSI target. It must contain 12 to 16 characters.

This password must differ from the password in the **Mutual Chap Secret** field.

Mutual Chap Secret

The mutual CHAP password. This password must be identical to the password configured on the iSCSI target. It must contain 12 to 16 characters.

This password must differ from the password in the **Chap Secret** field.

For more information, see the documentation " iSCSI Boot for PRIMERGY Server with Intel Network Controllers".

Boot Parameters for SAN boot in the case of Fibre Channel ports or functions

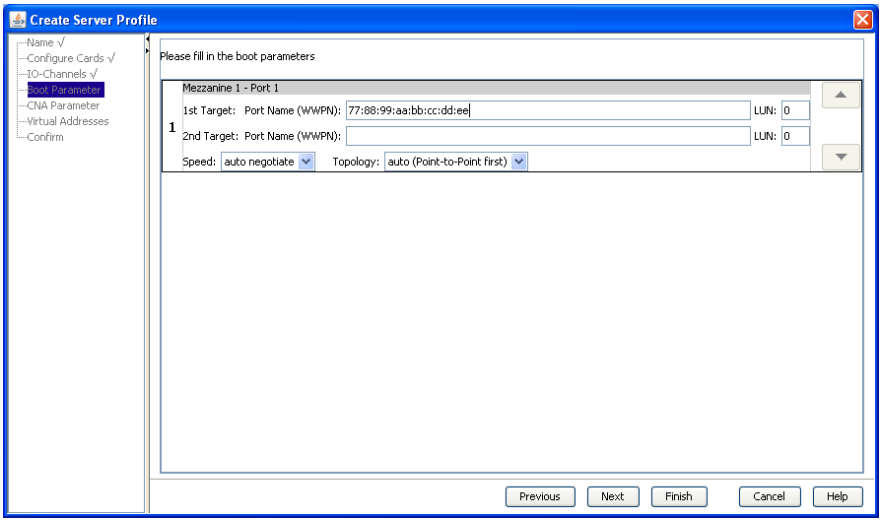

Figure 46: Boot Parameter step (SAN boot)

1st Target

In **1st Target**, you configure the boot device.

Port Name
WWPN (worldwide port name) of the port for the boot device.

LUN
LUN (logical unit number) address of the boot device. The default value for the field is 0.

2nd Target
In **2nd Target**, you configure the backup boot device (optional).

Port Name
WWPN (worldwide port name) of the port for the backup boot device.

LUN
LUN (logical unit number) address of the backup boot device. The default value for the field is 0.

Speed
Transmission speed used by the selected port. By default, **auto negotiate** is selected. If auto negotiation is not selected on the external switch, you configure the value using the drop-down menu according to the settings on the external switch. The following values are available:

auto negotiate
Default

The transmission speed is negotiated with the external switch.

1Gbit/s
1 Gbit/s full-duplex

2 Gbit/s
2 Gbit/s full-duplex

4 Gbit/s
4 Gbit/s full-duplex

8 Gbit/s
8 Gbit/s full-duplex

Topology
Type of port connection with the external SAN network. Possible values:

auto (loop first)

auto (Point-to-Point first)

Point-to-Point

Arbitrated loop

By default, **auto (Point-to-Point first)** is set.

5.4.3.5 CNA Parameter step (Create Server Profile wizard)

CNA Parameter is the next step in the **Create Server Profile** wizard. In this step, you specify the parameter for CNA functions.

This step is only shown if CNA onboard ports are selected or at least one CNA mezzanine card is specified in the second step of the **Create Server Profile** wizard.

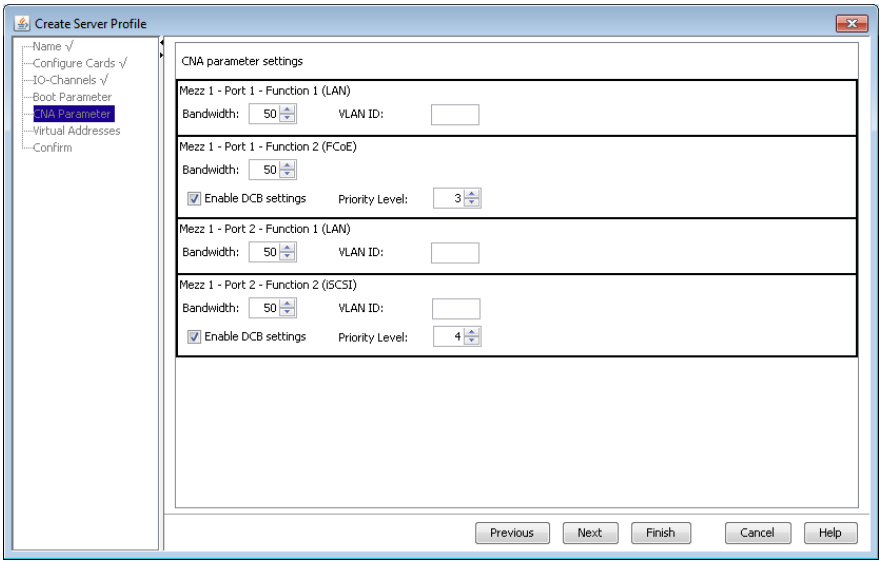

Figure 47: CNA Parameter step

Bandwidth

The share of the bandwidth in percent that is assigned to this function. If the sum of all bandwidths of one IO-channel is not 100, the values are internally adjusted accordingly.

This is the bandwidth reserved for the FCoE function. The FCoE function might share the complete bandwidth of 10 Gb with other functions. A value of 60, for example, means that a bandwidth of at least 6 Gb/sec is reserved for the FCoE packages.

Vlan ID

optional

The Vlan ID that is used by this function. This field is not available for FCoE functions.

Enable DCB Settings

This option is available only for FCoE and iSCSI functions. It enables the DCB (Data Center Bridging) feature of the connection blade. The DCB settings here are specific configuration settings in a DCB-enabled switching device.

This option should be enabled if the profile will be used in a blade server chassis where the corresponding physical CNA port will be connected to a PY CB Eth Switch/IBP 10 Gb 18/8 (SBAX2) connection blade in IBP mode.

If the port of the server blade, where the profile will be used, is connected to a LAN pass-thru module or a PY CB Eth Switch/IBP 10 Gb 18/8 (SBAX2) connection blade in switch mode, the option **Enable DCB settings** should not be set.

Priority Level

The priority level. Possible values are 0 to 7; for FCoE this is by default the value 3, for iSCSI the default value is 4.

5.4.3.6 Virtual Addresses step (Create Server Profile wizard)

Virtual Addresses is the next step in the **Create Server Profile** wizard. In this step, you specify the virtual addresses for each port. It is only shown if **Use virtual addresses** is activated in the **IO-Channels** step of the **Create Server Profile** wizard.

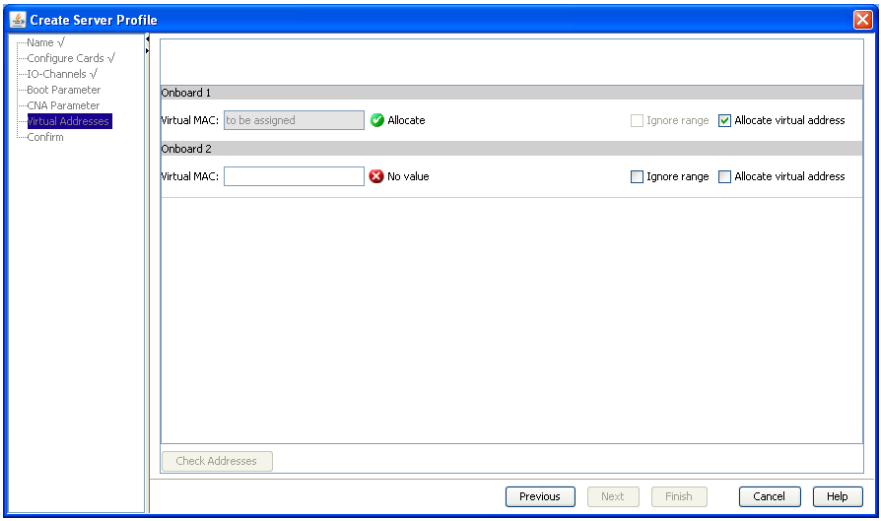

Figure 48: Virtual Addresses step

Virtual MAC, Virtual WWNN, Virtual WWPN, Virtual E-MAC

The virtual address. Virtual MAC and E-MAC (Enode MAC) addresses must have the structure **xx:xx:xx:xx:xx:xx** while virtual WWN addresses must have the structure **xx:xx:xx:xx:xx:xx:xx:xx**. Each **x** represents a hexadecimal character (0-9, a-f, A-F).

The Enode MAC address is the MAC address of an FCoE function.

Ignore range

If this option is selected, it is not checked whether the given virtual address is in the range that was specified when you installed VIOM.

Allocate virtual address

VIOM automatically assigns a virtual address after you exit the wizard.

Automatic assignment is only possible if you specified address ranges when you installed VIOM.

Next to each virtual address, you see the status of the address. The status can have the following values:

Status	Significance
Allocate (OK)	The address is automatically assigned.
OK (OK)	The address is valid and has not been assigned yet.
Not checked (warning)	The address entered has not been checked yet.
No value (error)	Either an address was not specified or **Allocate virtual address** was not selected.
Not unique (error)	The same address is in use for several ports.
Out of range (error)	The address is outside the specified address range.
Syntax error (error)	The address is syntactically incorrect.
Already used (error)	The address is already in use

Check Addresses

A check is performed to determine whether the addresses are currently in use in other profiles and whether they are within the range specified. This button is only active if for at least one address **Allocate virtual address** has not been selected and all virtual addresses have been entered correctly.

5.4.3.7 Confirm step (Create Server Profile wizard)

Confirm is the last step in the **Create Server Profile** wizard. In this step, you can check the entries you have made once again.

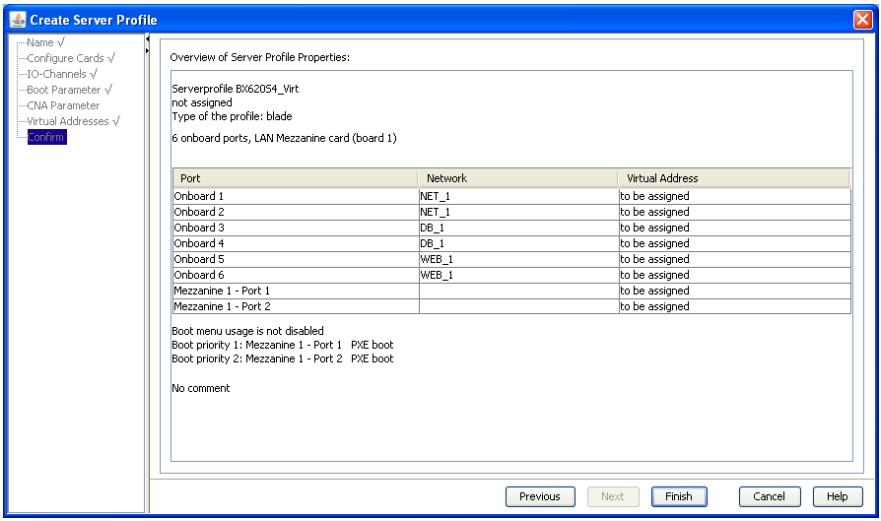

Figure 49: Confirm step

5.4.4 Edit Server Profile wizard

You use this wizard to modify a server profile.

The **Edit Server Profile** wizard comprises several dialog boxes to guide you through the individual steps. All required steps are displayed in the tree structure on the left.

To open the **Edit Server Profile** wizard, click the **Edit** button in the area on the right or select **Edit** in the context menu of the selected server profile.

5.4.4.1 Name step (Edit Server Profile wizard)

Name is the first step in the **Edit Server Profile** wizard. In the first step, you can modify the name, its intended target, and the description of the server profile.

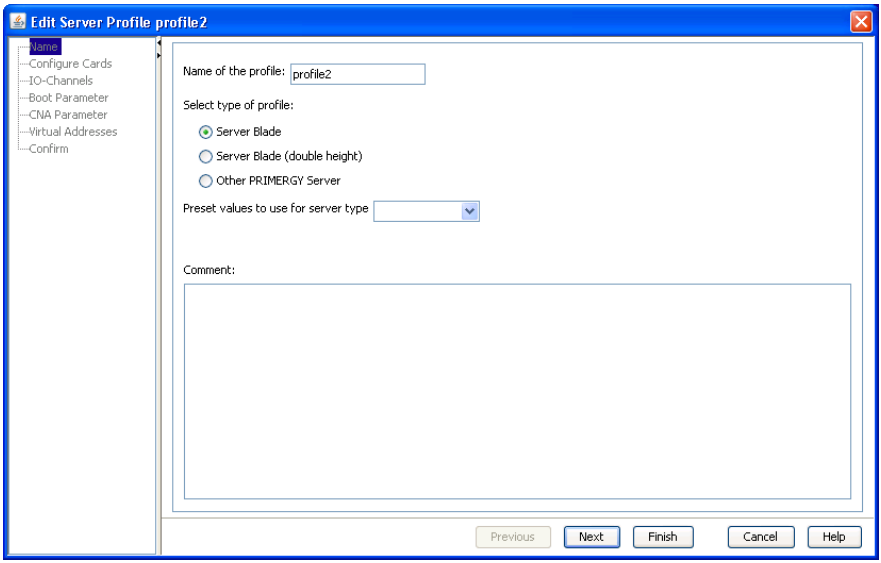

Figure 50: Name step

Name of the profile

Name of the server profile. If a profile already exists with this name or the name is invalid, the name is marked in red.

Select type of profile

The target type to which this profile should be assignable:

Server Blade

This profile can be assigned to slots of a blade server.

Server Blade (double height)

This profile can be assigned to two slots of a BX900 which are one over the other.

Other PRIMERGY Server

The profile can be assigned to a PRIMERGY rack server.

Preset values to use for server type

Server model (optional).

The number and type of LAN ports under **Onboard IO channels** are adjusted automatically according to the server model you select. The number of LAN ports and mezzanine/PCI cards cannot exceed the maximum possible value for the selected server model.

Comment

Comments on a more detailed description of the profile (optional)

5.4.4.2 Configure Cards step (Edit Server Profile wizard)

Configure Cards is the next step in the **Edit Server Profile** wizard.

In this step, you specify the number and type of mezzanine cards as well as the number of I/O channels for each card and for onboard.

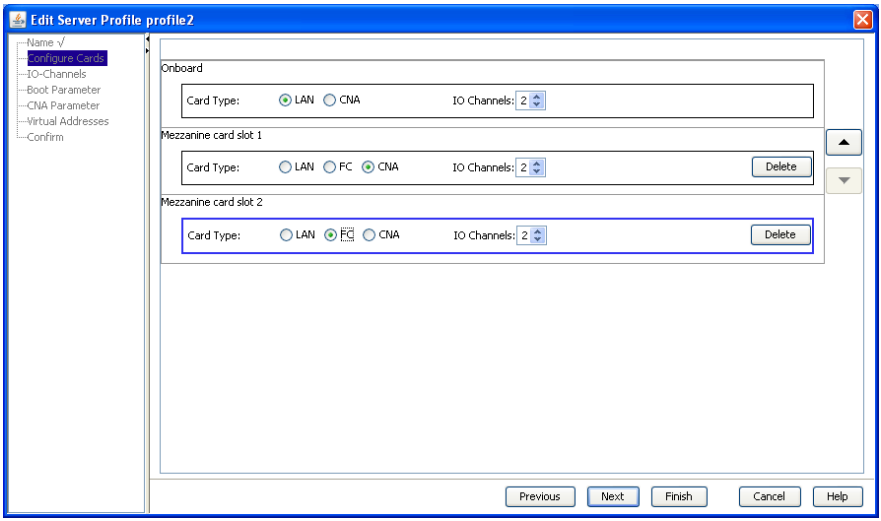

Figure 51: Configure Card step

In the **Onboard** box you can specify the type and number of onboard ports to be supported with this server profile. If you selected a server blade model under **Preset values to use for server type** in the **Name** step, the number and type of ports is initialized according to the model. You can adjust the number for the server profile. However, you cannot exceed the maximum possible value for the selected server model.

Depending on the selected profile and server type, there are a number of boxes, each representing a slot for a mezzanine or PCI card. If you want to configure a mezzanine or PCI card with this server profile, you must click the **Add** button in the corresponding slot to add a card. For each card you must specify its type:

LAN
LAN card. This card can have up to four ports.

FC
Fibre Channel card. This card can have up to two ports.

CNA
CNA card. This card can have up to two ports.

With **IO Channels** you specify the number of ports for the card to be supported with this server profile.

If you do not use an already configured mezzanine card, you can remove it with the **Delete** button.

With the arrow buttons on the right you can change the position of a mezzanine/PCI card. To do this, you must select a mezzanine/PCI card by clicking in its box (a blue border indicates that it is selected). You can use the up and down arrows to move the selected card by one position. If there is already a card in the target position, the cards swap their places. Any configuration for the mezzanine/PCI cards (e.g. boot configuration or virtual addresses) is retained.

5.4.4.3 IO-Channels step (Edit Server Profile wizard)

IO-Channels is the next step in the **Edit Server Profile** wizard.

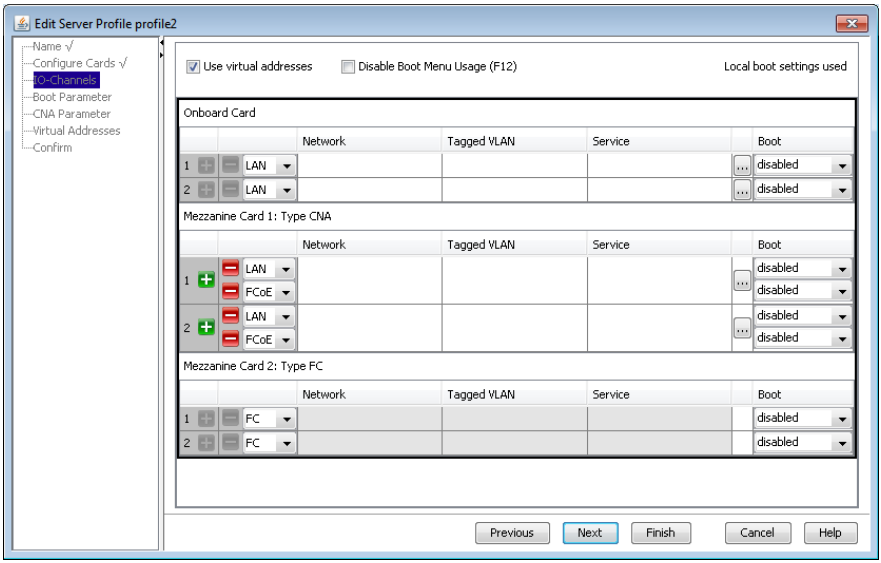

Figure 52: IO-Channels step

Local boot settings used

This message shows that the local boot settings are being used in the server profile. This message is not displayed if a boot device is used in the server profile.

Use virtual addresses

Uses virtual MAC addresses and WWN addresses with this profile. You can enter this information in another step within the wizard or VIOM can assign this information automatically.

Disable Boot Menu Usage (F12)

Prevents the VIOM boot settings from being overwritten on your local computer.

SMUX setting

This is only visible if a second mezzanine card of a blade server profile is defined as a LAN mezzanine card.

Defines the fabric to which the second mezzanine card is routed:

Fabric 3
All paths are routed to fabric 3.

Fabric 4
All paths are routed to fabric 4.

Fabric 3 & 4
LAN1 is routed to fabric 3 and LAN2 to fabric 4.

For more information on SMUX settings, see the documentation "PRIMERGY BX900 Blade Server Systems - ServerView Management Blade S1".

The upper table displays the onboard ports (up to 6). Another table is displayed for each available mezzanine/PCI card configured in the previous step. The tables have the following columns:

Column	Significance
first column	Port number of the onboard or the mezzanine/PCI card port. On CNA cards the type of the function (LAN, FCoE, iSCSI) can be selected from the drop-down list. Another function can be added to this port with the  button. Functions can be removed with the  button.

Column	Significance
Network	Name of the network. You can specify a network for each LAN or CNA port, but not for profiles of type **Rack Server**. If you want to use the profile on blade servers with IBP modules, you can specify a network for each LAN or CNA port. If you work with blade servers that have non-VIOM-capable LAN modules (Open Fabric mode), do not specify a network as it is not possible to define networks on these modules. To enter a network name, click the table cell to switch to edit mode. You can also open a network selection dialog box via the "..." button. In it you can select a managed blade server chassis from the selection list. The networks defined for this chassis that are suitable for the selected port are then displayed. To select a network, double-click the name or select the name and click the **Add** button. Make sure that the networks entered here are/will be configured before the profile is activated on the corresponding blade server. As long as a network does not yet exist, the server profile can be created with this network, but cannot yet be assigned to a slot. If you wish to exit the network selection without selecting a network, click the **Close** button or another input field.
Tagged VLAN	Names of tagged VLAN networks. You can specify tagged VLAN networks for each LAN or CNA port, but not for profiles of type **Rack Server**. If you specify more than one tagged VLAN for a port, the names must be separated by commas. If you use the network selection box, the name of the chosen network is added to the **Tagged VLAN** column if you use the **Add tagged** button.

Column	Significance
Service	Names of the service networks. You can specify service networks for each LAN or CNA port, but not for profiles of type **Rack Server**. If you specify more than one service network for a port, the names must be separated by commas. If you use the network selection box, the name of the chosen network is replaced in the **Network** column or added to the **Service** column depending on the type of the selected network.
Boot	To configure the port as the boot device, select **PXE boot**, **iSCSI boot**, or **SAN boot** from the selection list. If you select **disabled**, this port is not a boot device. You can define up to four boot devices in a profile. The values available in the selection list depends on type of port or function. If you configure an iSCSI boot device or SAN boot device, you must specify additional boot settings in the **Boot Parameter** step.

5.4.4.4 Boot Parameter step (Edit Server Profile wizard)

Boot Parameter is the next step in the **Edit Server Profile** wizard. In this step, you can modify the boot device configuration and the boot parameter order. This step is only shown if you configured a iSCSI boot or SAN boot or at least two channels as boot devices in the previous step of the **Edit Server Profile** wizard.

If you configured at least two channels as boot devices in the previous step, use the arrow up and arrow down buttons to change the boot order.

The selection you made at the second step (iSCSI boot or SAN boot) determines which fields are displayed in this dialog box.

Boot Parameters for an iSCSI boot (LAN ports) or iSCSI function on CNA mezzanine card

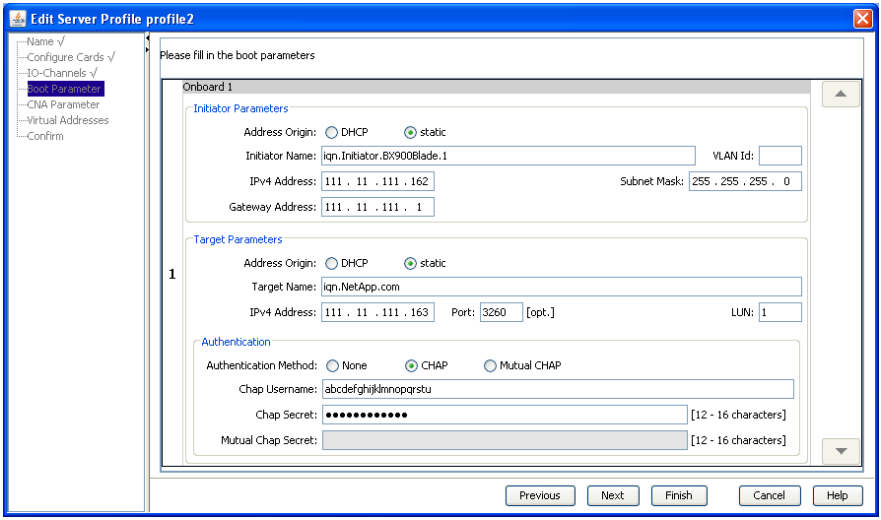

Figure 53: Boot Parameter step (iSCSI boot)

Initiator Parameters

Address Origin

DHCP

The system tries (in the case of an iSCSI boot) to obtain the client IP address, subnet mask, and gateway IP address from a DHCP server. Only the initiator name and (optional) the VLAN ID must be specified here.

static

A static client IP address, subnet mask, and gateway IP address must be specified.

Initiator Name

Name of the iSCSI initiator to be used (in the case of an iSCSI boot) for the connection to the iSCSI target.

IPv4 address

Static client IP address to be used for this port. The port will use this IP address for the entire iSCSI session. You can enter an IP address in this field if **Address Origin: static** is selected.

Subnet Mask

IP subnet mask. This should be the IP subnet mask of the network used to connect this port (in the case of an iSCSI boot). You can enter a subnet mask in this field if **Address Origin: static** is selected.

Gateway Address

IP address of the network gateway. You can enter the gateway address in this field if **Address Origin: static** is selected. This is necessary if the iSCSI target is in a subnetwork other than the subnetwork of the selected iSCSI boot port.

If iSCSI initiator and target are in the same network segment, no gateway is needed. The gateway address should be set to the value **0.0.0.0**.

Target Parameters

Address Origin

DHCP

The system tries (in the case of an iSCSI boot) to obtain the name of the iSCSI target, the IP address of the iSCSI target, the IP port number, and the SCSI LUN ID from a DHCP server in the network.

static

Static name for the iSCSI target, a static IP address for the iSCSI target, a static IP port number, and a static SCSI LUN ID.

Target Name

IQN name of the iSCSI target. You can enter a name in this field if **Address Origin: static** is selected.

IPv4 address
IP address of the iSCSI target. You can enter an IP address in this field if **Address Origin: static** is selected.

Port
TCP port number (default: 3260 for iSCSI). You can enter a port number in this field if **Address Origin: static** is selected. (optional)

LUN
LUN ID of the boot disk on the SCSI target. You can enter a LUN ID in this field if **Address Origin: static** is selected.

Authentication

Authentication Method

None
No authentication is used.

CHAP
CHAP authentication is activated for this port. CHAP allows the target to authenticate the initiator. After activating CHAP, you must enter a user name and password for the target.

Mutual CHAP
Mutual CHAP authentication is activated for this port. Mutual CHAP allows the initiator to authenticate the target. After activating mutual CHAP authentication, you must enter a user name, a password for the target, and a mutual CHAP password.

Chap Username
CHAP user name. The name must be identical to the name configured on the iSCSI target.

Chap Secret
CHAP password. This password must be identical to the password configured on the iSCSI target. It must contain 12 to 16 characters.

This password must differ from the password in the **Mutual Chap Secret** field.

Mutual Chap Secret

The mutual CHAP password. This password must be identical to the password configured on the iSCSI target. It must contain 12 to 16 characters.

This password must differ from the password in the **Chap Secret** field.

For more information, see the documentation " iSCSI Boot for PRIMERGY Server with Intel Network Controllers".

Boot Parameters for SAN boot in the case of Fibre Channel ports or functions

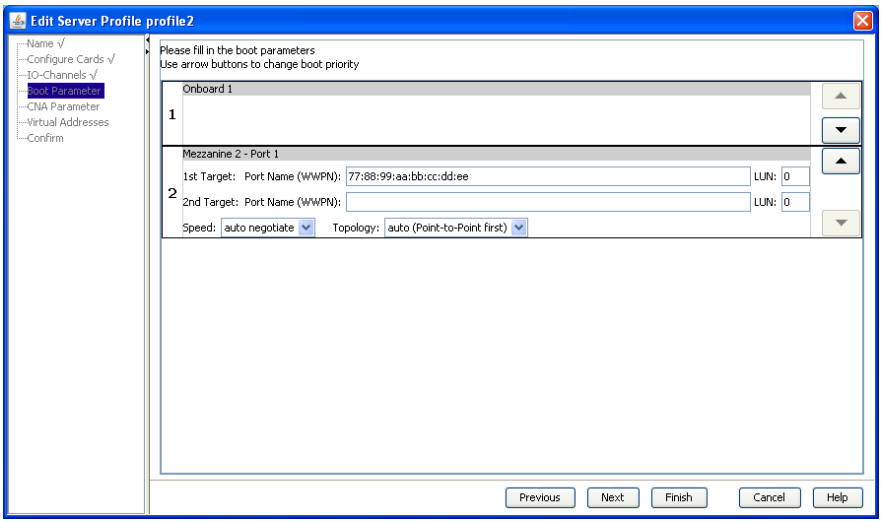

Figure 54: Boot Parameter step (SAN boot)

1st Target

In **1st Target**, you configure the boot device.

Port Name

WWPN (worldwide port name) of the port for the boot device.

LUN

LUN (logical unit number) address of the boot device. The default value for the field is 0.

2nd Target

In **2nd Target**, you configure the backup boot device.

Port Name

WWPN (worldwide port name) of the port for the backup boot device.

LUN

LUN (logical unit number) address of the backup boot device. The default value for the field is 0.

Speed

Transmission speed used by the selected port. By default, **auto negotiate** is selected. If auto negotiation is not selected on the external switch, you configure the value using the drop-down menu according to the settings on the external switch. The following values are available:

auto negotiate

Default

The transmission speed is negotiated with the external switch.

1Gbit/s

1 Gbit/s full-duplex

2 Gbit/s

2 Gbit/s full-duplex

4 Gbit/s

4 Gbit/s full-duplex

8 Gbit/s

8 Gbit/s full-duplex

Topology

Type of port connection with the external SAN network. Possible values:

auto (loop first)

auto (Point-to-Point first)

Point-to-Point

Arbitrated loop

By default, **auto (Point-to-Point first)** is set.

5.4.4.5 CNA Parameter step (Edit Server Profile wizard)

CNA Parameter is the next step in the **Edit Server Profile** wizard. In this step, you specify the parameter for CNA functions.

This step is only shown if at least one FCoE function for a CNA mezzanine card is specified in the second step of the **Edit Server Profile** wizard.

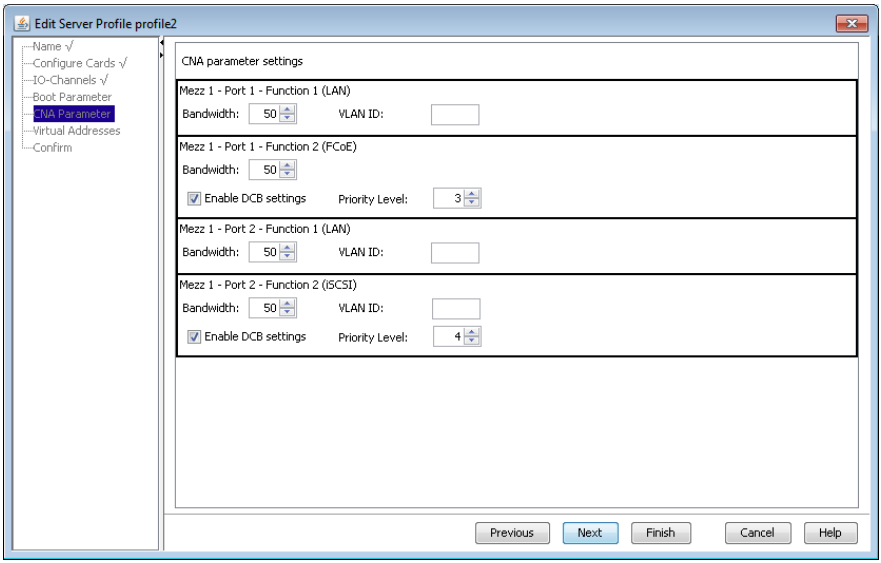

Figure 55: CNA Parameter step

Enable DCB Settings

This option is available only for FCoE and iSCSI functions. It enables the DCB (Data Center Bridging) feature of the connection blade. The DCB settings here are specific configuration settings in a DCB-enabled switching device.

This option should be enabled if the profile will be used in a blade server chassis where the corresponding physical CNA port will be connected to a PY CB Eth Switch/IBP 10 Gb 18/8 (SBAX2) connection blade in IBP mode.

If the port of the server blade, where the profile will be used, is connected to a LAN pass-thru module or a PY CB Eth Switch/IBP 10 Gb 18/8 (SBAX2) connection blade in switch mode, the option **Enable DCB settings** should not be set.

Priority Level

The priority level. Possible values are 0 to 7; for FCoE this is by default the value 3, for iSCSI the default value is 4.

Bandwidth

The share of the bandwidth in percent that is assigned to this function. If the sum of all bandwidths of one IO-channel is not 100, the values are internally adjusted accordingly.

This is the bandwidth reserved for the FCoE function. The FCoE function might share the complete bandwidth of 10 Gb with other functions. A value of 60, for example, means that a bandwidth of at least 6 Gb/sec is reserved for the FCoE packages.

5.4.4.6 Virtual Addresses step (Edit Server Profile wizard)

Virtual Addresses is the next step in the **Edit Server Profile** wizard. In this step, you specify the virtual addresses for each port. It is only shown if **Use virtual addresses** is activated in the second step of the **Edit Server Profile** wizard.

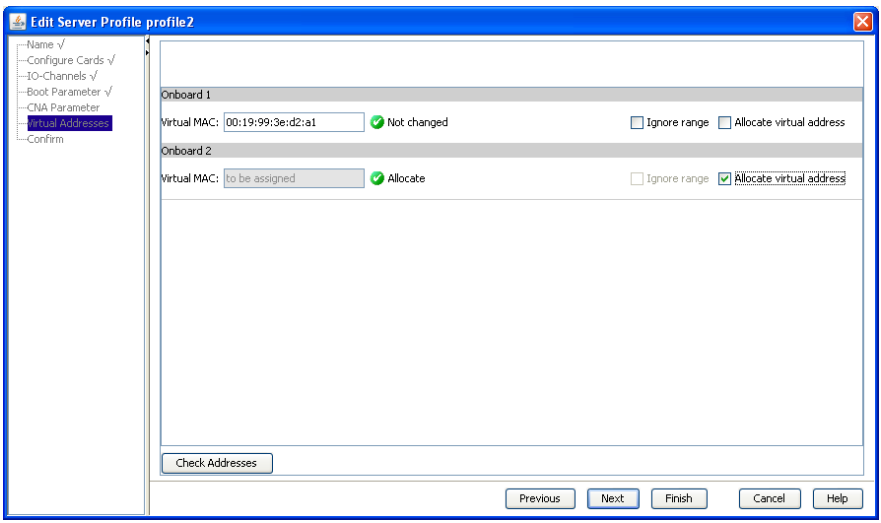

Figure 56: Virtual Addresses step

Virtual MAC, Virtual WWNN, Virtual WWPN, Virtual E-MAC

The virtual address. Virtual MAC and E-MAC (Enode MAC) addresses must have the structure **xx:xx:xx:xx:xx:xx** while virtual WWN addresses must have the structure **xx:xx:xx:xx:xx:xx:xx:xx**. Each **x** represents a hexadecimal character (0-9, a-f, A-F).

The Enode MAC address is the MAC address of an FCoE function.

Ignore range

If this option is selected, it is not checked whether the given virtual address is in the range that was specified when you installed VIOM.

Allocate virtual address

VIOM automatically assigns a virtual address after you exit the wizard.

Automatic assignment is only possible if you specified address ranges when you installed VIOM.

Next to each virtual address, you see the status of the address. The status can have the following values:

Status	Significance
Allocate (OK)	The address is automatically assigned.
Not changed (OK)	The address is not changed.
OK (OK)	The address is valid and has not been assigned yet.
Not checked (warning)	The address entered has not been checked yet.
No value (error)	Either an address was not specified or **Allocate virtual address** was not selected.
Not unique (error)	The same address is in use for several ports.

Status	Significance
Out of range (error)	The address is outside the specified address range.
Syntax error (error)	The address is syntactically incorrect.
Already used (error)	The address is already in use

Check Addresses

A check is performed to determine whether the addresses are currently in use in other profiles and whether they are within the range specified. This button is only active if for at least one address **Allocate virtual address** has not been selected and all virtual addresses have been entered correctly.

5.4.4.7 Confirm step (Edit Server Profile wizard)

Confirm is the next step in the **Edit Server Profile** wizard. In this step, you can check the entries you have made once again.

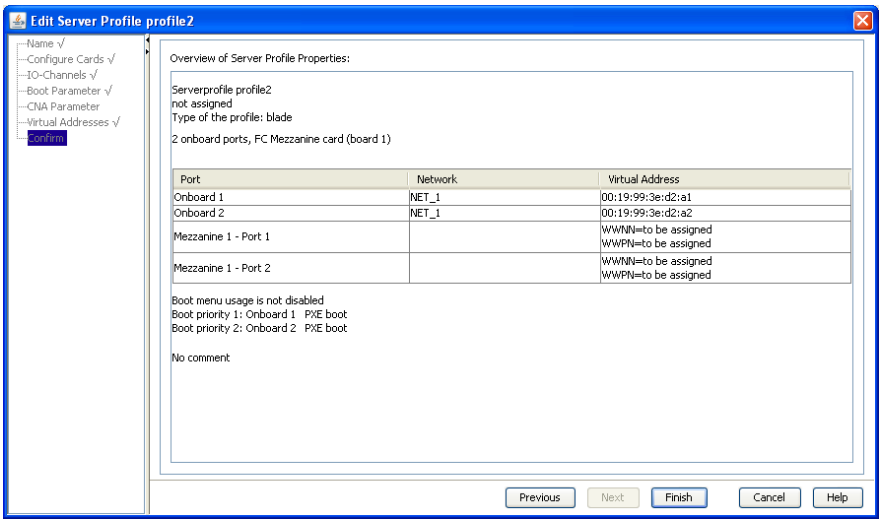

Figure 57: Confirm step

5.4.5 Save Configuration wizard

You use the **Save Configuration** wizard to save and to restore backup files as well as to delete them on the management station.

The **Save Configuration** wizard comprises several dialog boxes to guide you through the individual steps. All required steps are displayed in the tree structure on the left.

You launch the wizard using the **Configuration Backup / Restore** button on the **Virtual-IO Manager** tab.

5.4.5.1 Select Action step (Configuration Backup/Restore wizard)

You save and restore backup files as well as delete them on the management station using a wizard. You launch the wizard using the **Configuration Backup / Restore** button on the **Virtual-IO Manager** tab.

Select Action is the first step in the Configuration Backup/Restore wizard. In this step you select the action that you want to carry out.

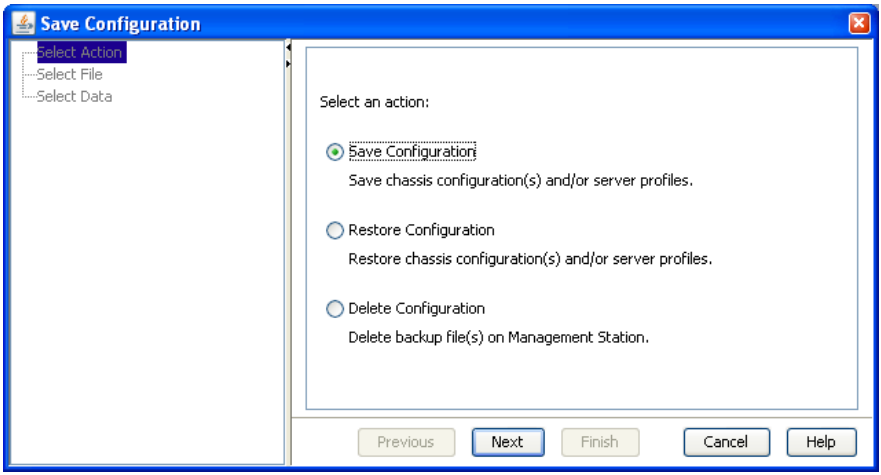

Figure 58: Select Action step

Save Configuration

Select **Save Configuration** to save a configuration in a file.

Restore Configuration

Select **Restore Configuration** to restore a configuration from a file.

Delete Configuration

Select **Delete Configuration** to delete backup files you no longer need on the central management station.

You delete backup files saved locally using the means available on the operating system of the local computer

5.4.5.2 Select File step (Save Configuration Wizard)

Select File is the second step in the **Save Configuration** wizard. In this step you select the computer on which you want to save the configuration.

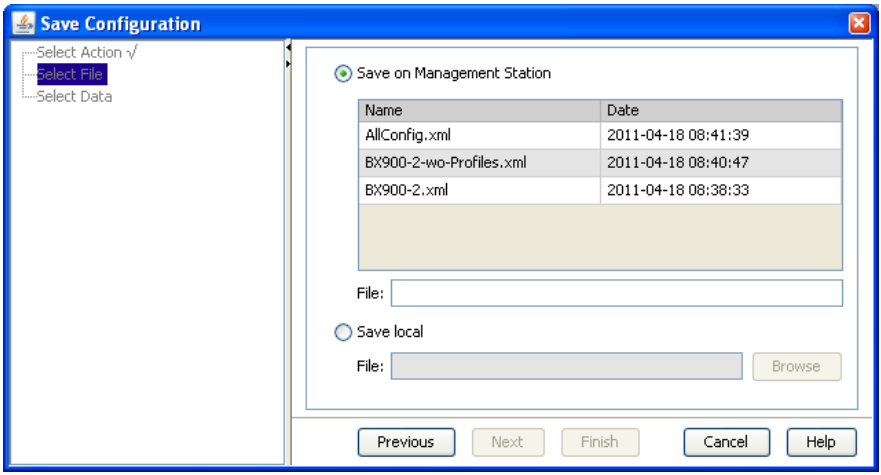

Figure 59: Select File step (Save Configuration wizard)

Save on Management Station

Saves the configuration on the management station.

Column	Significance
Name	Name of existing backup files that will be overwritten

Column	Significance
Date	Date and time when the backup file is created.

You can sort the list of existing backup files according to file name or date by clicking the table headline accordingly.

File

Name for the backup file. You can also create a file in subdirectories by specifying the entire path name (e.g. **directory/file**). If the required directories do not yet exist, they are created automatically. The backup files are assigned the **.xml** suffix automatically if you have not specified one.

Save local

Saves the configuration in a file on the computer on which the GUI runs.

File

Name for the backup file. It is strongly recommended to specify a complete path (on Windows including drive letter) because otherwise the file location depends on the used browser and operating system.

Browse

Opens the file selection dialog box in which you can navigate to the desired folder and then select an existing backup file or specify the name of the backup file.

If the file already exists, it will be overwritten.

5.4.5.3 Select File step (Restore Configuration wizard)

Select File is the second step in the **Restore Configuration** wizard. In this step you select the the backup file from which you want to restore the configuration.

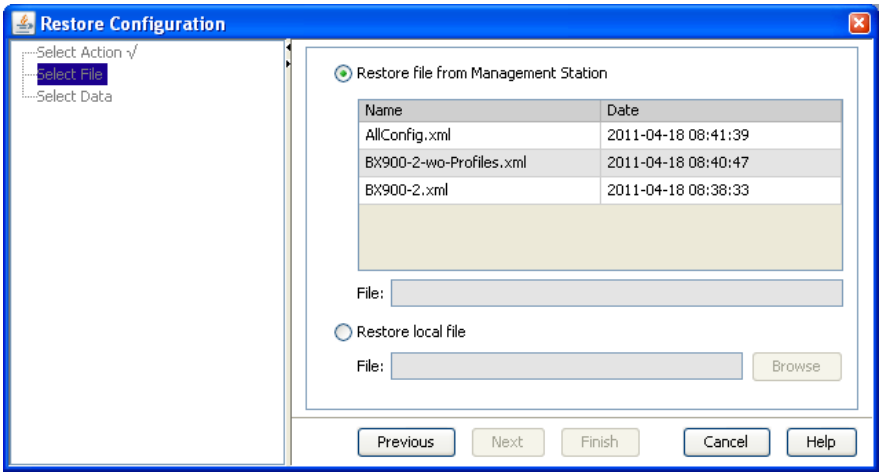

Figure 60: Select File step (Restore Configuration wizard)

Restore file from Management Station

Backup file from the management station.

Column	Significance
Name	Name of an existing backup file.
Date	Date and time when the backup file is created.

You can sort the list of existing backup files according to file name or date by clicking the table headline accordingly.

File

Name for the backup file.

Restore local file

A backup file on the local computer on which the GUI runs

File

Name for the backup file.

Browse

Opens the file selection dialog box in which you can select the relevant backup file.

5.4.5.4 Select File step (Delete Backup Files wizard)

Select File is the second step in the **Delete Backup Files** wizard. In this step you select the backup files you want to delete on the management station.

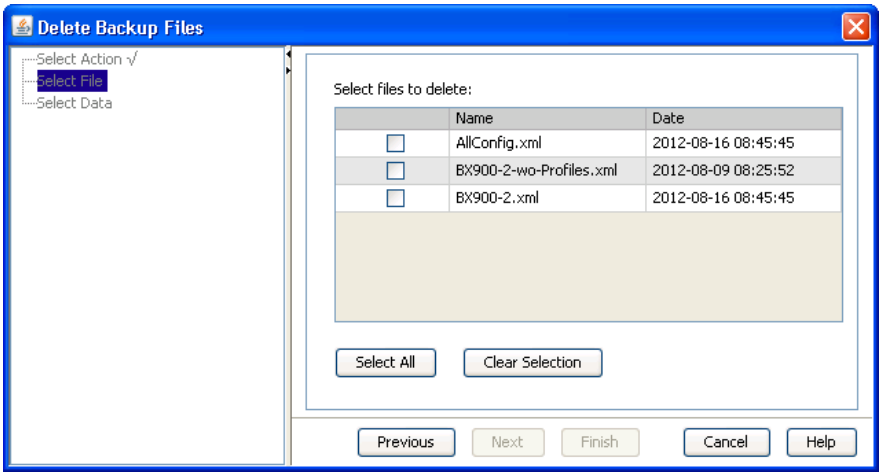

Figure 61: Select File step (Delete Backup Files wizard)

The table displays the backup files. You can sort the list of backup files according to file name or date by clicking the table headline accordingly.

Column	Significance
	Selection box
Name	Name of the backup file
Date	Date and time, when the backup file is created

Select All

Selects all the backup files.

Clear Selection

Deselects all the backup files.

5.4.5.5 Select Data step (Save Configuration wizard)

Select Data is the third step in the **Save Configuration** wizard. In this step you select the servers whose configuration you wish to save.

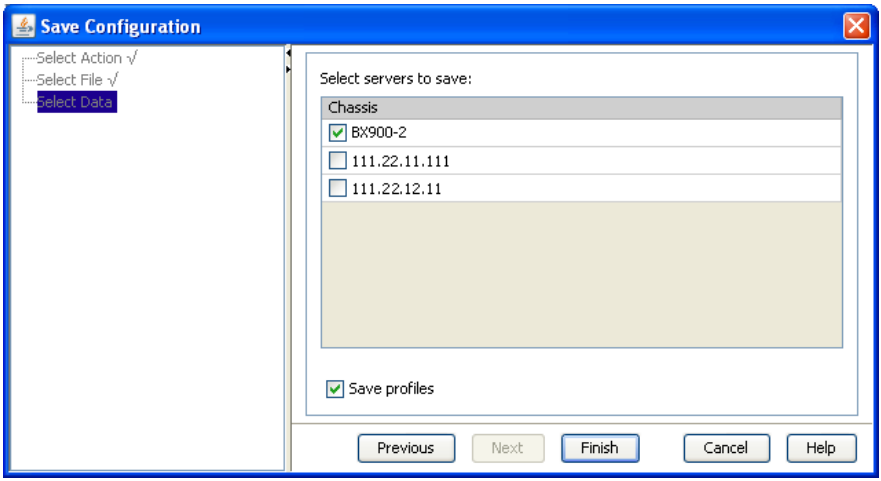

Figure 62: Select Data step (Save Configuration wizard)

Select servers to save

Servers whose configuration you wish to save. Do not select any servers if you only wish to save profiles.

Save Profiles

Specify whether server profiles are to be saved too.

It is essential that you save the profiles along with the configuration if the assigned profiles are to be reassigned once the configuration has been restored.

5.4.5.6 Select Data step (Restore Configuration wizard)

Select Data is the third step in the **Restore Configuration** wizard. In this step you define the data you wish to restore.

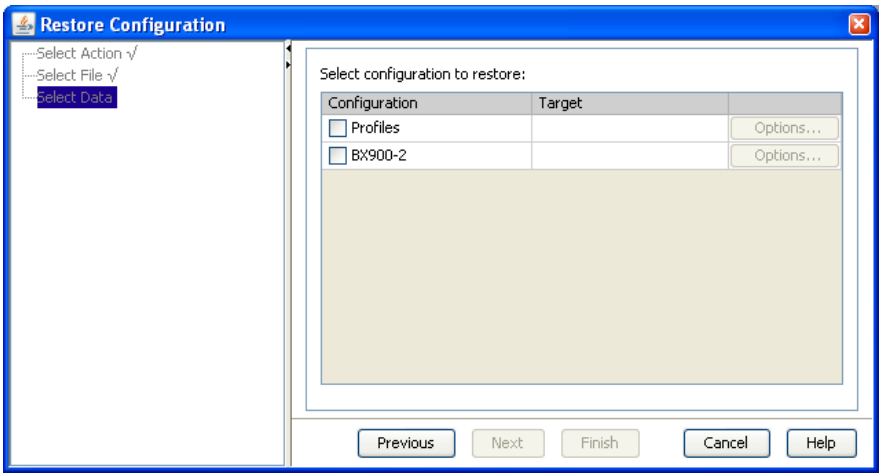

Figure 63: Select Data step (Restore Configuration wizard)

Select the configuration to restore

In the **Configuration** column you select the configuration(s) to restore.

Column	Significance
Configuration	Selection box and configurations that could be restored.
Target	For blade servers only: The target to which the configuration will be restored is specified.
	Options... button. It opens the **Restore Options** dialog box.

You can select following configurations to restore:

Profiles

Restore the server profiles.

The **Restore Options** dialog box opens automatically in which you specify additional parameters (see section "Restore Options dialog box (server profiles)" on page 243).

<server>

Restore the configuration saved for the selected server.

For blade servers, the **Restore Options** dialog box opens automatically in which you specify additional parameters (see section "Restore Options dialog box (servers)" on page 241) .

In addition, you can click the **Options...** button in the last column to open the **Restore Options** dialog box (see "Restore Options dialog box (servers)" on page 241 for each configuration.

For PRIMERGY rack servers, the three options in the **Restore Options** dialog box are disabled. In addition, **Reassign Profiles** is selected.

5.4.5.7 Select Data step (Delete Backup Files wizard)

Select Data is the third step in the **Delete Backup Files Configuration** wizard.

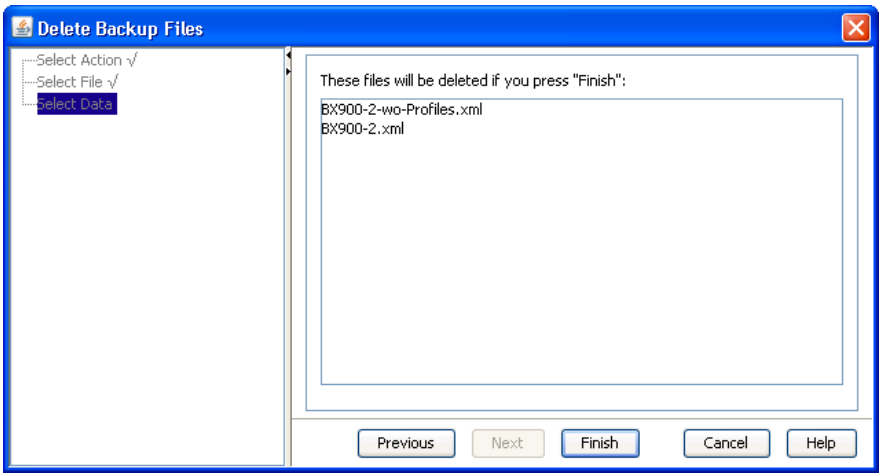

Figure 64: Select Data step (Delete Backup Files wizard)

In this step you will see the names of the files you wish to delete. The files are deleted after you have clicked **Finish**.

5.5 Dialog boxes

Virtual-IO Manager provides the following dialog boxes:

- **Authentication** dialog box
- **Licenses Information** dialog box
- **Restore Options** dialog box
- **Select Profile** dialog box

5.5.1 Authentication dialog box (single blade server)

The **Authentication** dialog box is displayed when you click the **Manage** or **Authentication** button on the **Setup** tab. This dialog box varies depending on what you have selected in the table on the **Setup** tab.

If you select a single blade server in this table, the following **Authentication** dialog box is shown.

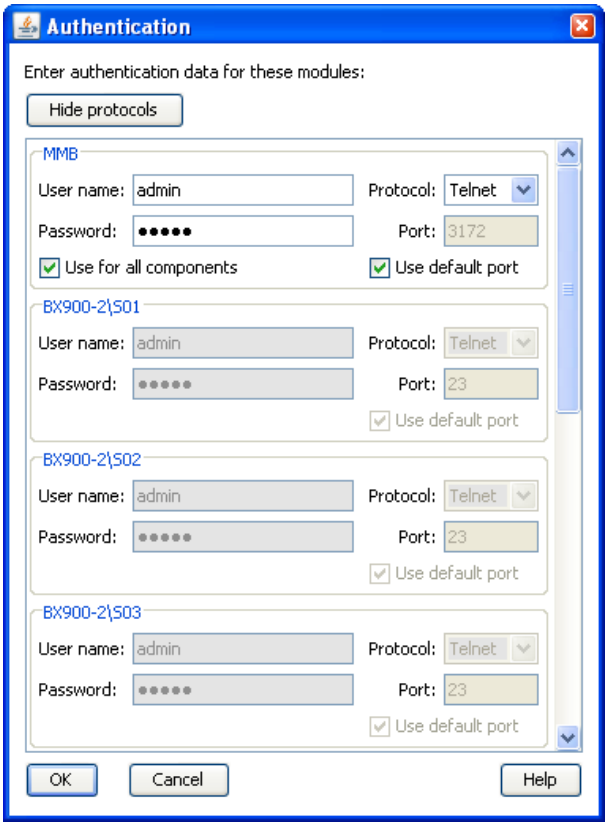

Figure 65: VIOM Manager authentication (single blade server)

In this dialog box, you enter the user names/passwords for the management blade and for the IBP modules, which VIOM can then use to access these modules.

For the MMB and each IBP, you must specify the user name and password with which VIOM can access the component. To make it easier for you, in the MMB area there is a checkbox called **Use for all components** which, when checked, sets the values for all IBPs to the ones specified for the MMB. So, if the user name and password are the same for all components (this should be the standard), you only need to specify them once.

For security reasons you should change the default user name and password combination. You can have the same user name and password for all components or different ones for each. You must specify this accordingly in this dialog box.

With the **Configure protocols** button you can expand the dialog box, so that you can also specifiy the protocol and port to be used by VIOM to access the components. These values should not be changed in a standard configuration.

User name
A valid user name with access rights to the MMB or IBP

Password
Password of the user ID

Use for all components
Sets the values for all IBPs to the ones specified for the MMB.

Configure protocols
Expands the dialog box, so that you can also specifiy the protocol and the port to be used by VIOM to access the components. These values should not be changed in a standard configuration.

Protocol
Select the protocol to be used for communication with the management blade or the I/O connection blades.

Port
If **Use default port** is not checked, you have the option of specifying other port numbers for the modules in this field.

Use default port
Uses the default port. The default port is dependent on the protocol and the relevant module.

For a management blade: Port 3172 for Telnet and Port 22 for SSH

For an I/O connection blade: Port 23 for Telnet and Port 22 for SSH

The **Port** input field is inactive if this option is checked.

For a standard configuration, these values should not be changed.

5.5.2 Authentication dialog box (PRIMERGY rack server)

The **Authentication** dialog box is displayed when you click the **Manage** or **Authentication** button on the **Setup** tab. This dialog box varies depending on what you have selected in the table on the **Setup** tab.

If you select one or more PRIMERGY rack server(s) in this table on the **Setup** tab, the following **Authentication** dialog box is shown.

Figure 66: VIOM Manager authentication (PRIMERGY rack server)

For each PRIMERGY rack server, you must to specify the user name and password for the iRMC with which VIOM can access this server. If more than one PRIMERGY rack server is selected in the table, there is a check-box called **Use for all components** which, when checked, sets the values for all PRIMERGY rack servers to the ones specified for the first PRIMERGY rack server. So, if the user name and password are the same for all PRIMERGY rack servers, you only need to specify them once.

For security reasons you should change the default user name and password combination. You can have the same user name and password for all com-

ponents or different ones for each. You must specify this accordingly in this dialog box.

With the **Configure protocols** button you can expand the dialog box, so that you can also specifiy the port to be used by VIOM to access the components. This value should not be changed in a standard configuration.

User name
A valid user name with access rights to the iRMC

Password
Password of the user ID

Trap Destination
IP address (IPv4 or IPv6) of the VIOM management station to which traps are sent in the case of connecting the PRIMERGY rack server to power. This field is preset with one IP address of the VIOM management station. It should only be changed if the management station is reachable by different IP addresses and another one should be used.

Use for all components
Sets the values for all PRIMERGY rack servers to the ones specified for first PRIMERGY rack server.

Configure protocols
Expands the dialog box, so that you can also specifiy the protocol and the port to be used by VIOM to access the PRIMERGY rack server(s). These values should not be changed in a standard configuration.

Protocol
Select the protocol to be used for communication with the PRIMERGY rack servers.

Port
If **Use default port** is not checked, you have the option of specifying other port numbers for the PRIMERGY rack servers in this field.

Use default port

Uses the default port. The default port is dependent on the protocol and the relevant module.

For a PRIMERGY rack server: Port 623

The **Port** input field is inactive if this option is checked.

For a standard configuration, these values should not be changed.

5.5.3 Authentication dialog box (PRIMERGY rack server and blade server)

The **Authentication** dialog box is displayed when you click the **Manage** or **Authentication** button on the **Setup** tab. This dialog box varies depending on what you have selected in the table on the **Setup** tab.

If you select several servers with at least one blade server in this table, you only specify one user name and password. These will be used for all MMBs and IBPs of all selected blade servers and for all selected PRIMERGY rack servers. For protocol, port, and trap destination, the default values are used. If this situation is not suitable for you, please select fewer servers in the table before you click the **Manage** or **Authentication** button.

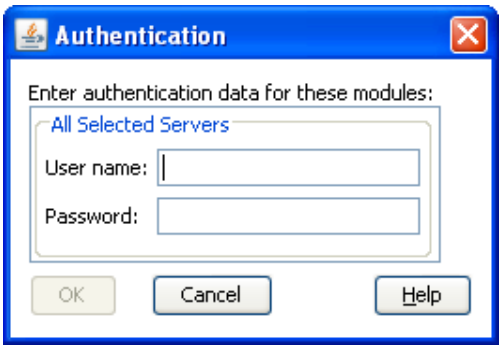

Figure 67: VIOM Manager authentication (PRIMERGY rack server and blade server)

User name
A valid user name with access rights to the MMBs, IBPs, and iRMCs

Password
Password of the user ID

5.5.4 Licenses Information dialog box

The **Licenses Information** dialog box displayes a table showing the relevant information of the license.

This dialog box opens when you click the **Show Licenses** button on the **Virtual-IO Manager** tab.

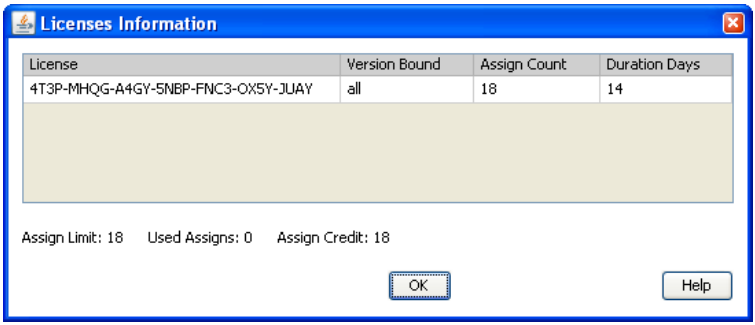

Figure 68: Licenses Information

In the table there is one row per license:

Column	Significance
License	License code
Version Bound	VIOM version if the license is version bound (currently not yet supported)
Assign Count	Number of server profiles that this license permits you to assign using the Virtual-IO Manager
Duration Days	Validity period (only applies to demo licenses)

Assign Limit
Maximum number of server profiles that can be assigned with the licenses

Used Assigns
Number of server profiles currently being assigned

Assign Credit
Number of server profiles that can still be assigned with the licenses

OK
Closes the **Licenses Information** dialog box.

5.5.5 Preferences dialog box

The **Preferences** dialog box allows you to set user preferences.

This dialog box opens when you click the **Preferences** button on the **Virtual-IO Manager** tab.

In the **Preferences** dialog box are two tabs to set user preferences:

- **Display** tab
- **Trace** tab

Display tab

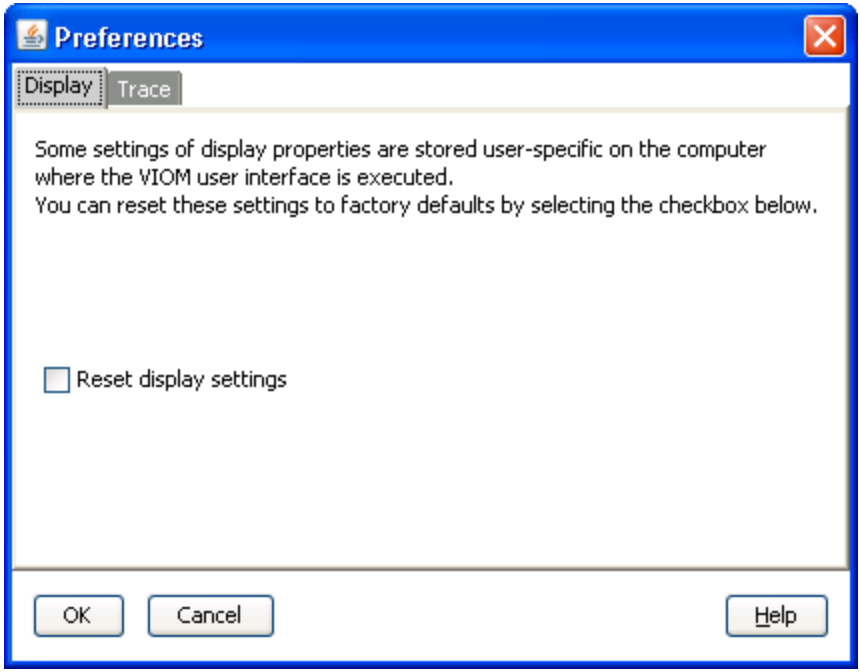

Figure 69: Display tab (Preferences dialog box)

The **Display** tab allows you to reset the display settings. If you change some display properties (e.g. the width of columns in a table), these are stored user-specific. These settings can be reset to factory defaults by selecting the **Reset display settings** check box.

Trace tab

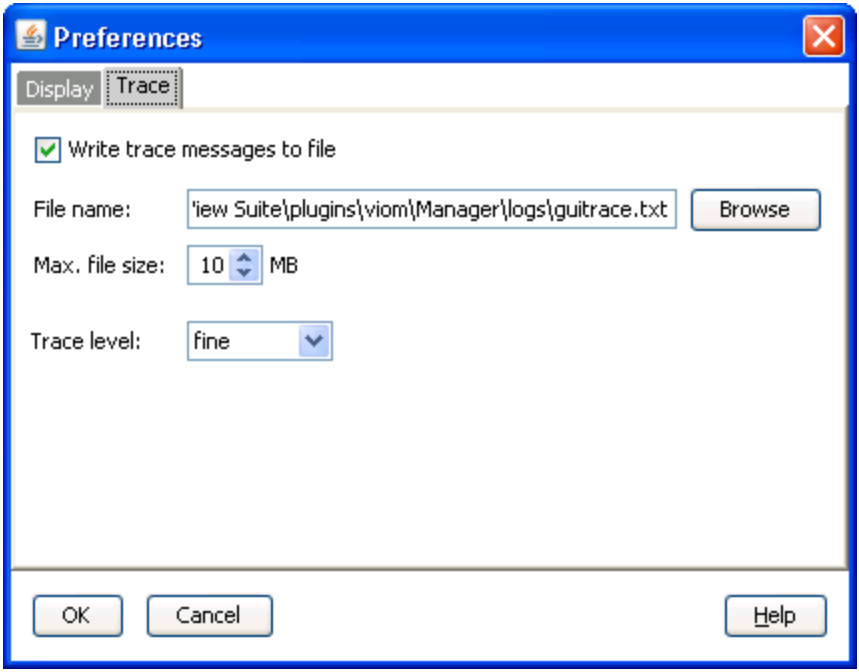

Figure 70: Trace tab (Preferences dialog box)

On the **Trace** tab, parameters for tracing the VIOM-GUI can be set. This is useful to create diagnostic data (see section "Creating diagnostic data" on page 356). On this tab you find the following options and input fields:

Write trace messages to file

Select this option to write trace output to file(s). If this option is not selected, the trace output is written to the Java Console only.

File name

Enter here the file name for the trace output. You can use the **Browse** button to select the file.

Max. file size

Specifies the maximal size of a trace file.

If the maximum file size is reached, the trace file will be renamed and a new one will be used. The renamed trace files will have a number appended, up to a maximum of ten possible back-up trace files. Therefore the trace files can use up to ten times the specified value of disk space.

Trace level

Specifies the amount of trace information. To create diagnostic data, **fine** is an appropriate value.

no tracing

Tracing is switched off.

severe

Only the most important information is traced.

info

Also some information about minor errors is traced.

fine

Information necessary for error diagnosis is traced.

call trace

More detailed information as **fine**.

finer

Even more information than in **call trace**.

data

Most detailed trace; currently not used.

5.5.6 Restore Options dialog box (servers)

The **Restore Options** dialog box opens when you select a blade server in the **Select Data** step of the **Restore Configuration** wizard. It also opens when you click the **Options...** button in the table row. In this dialog box you can specify additional parameters to restore blade server configurations.

For PRIMERGY rack servers, the **Restore Options** dialog box can also be opened by clicking the **Options...** button. But the three options in the **Restore Options** dialog box are disabled. In addition, **Reassign Profiles** is selected. So in this case, you cannot change any of the options.

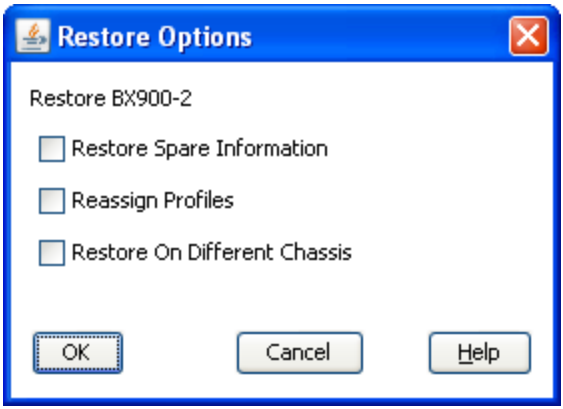

Figure 71: Restore Options dialog box (blade server)

Restore Spare Information

Specifies whether the information on the spare slots is restored.

Reassign Profiles

Specifies whether the profiles which were assigned when the backup was performed are reassigned. This option is only enabled if the backup contains profiles.

If **Profiles** is not selected in the **Select Data** step of the **Restore Configuration** wizard, it is selected automatically when you close this dialog box with **OK**. In this case, the **Restore only reassigned profiles** option (see "Restore Options dialog box (server profiles)" on page 243) will be selected for profiles.

Restore On Different Chassis

Specifies that the backup is restored on another blade server. If you select this option, a **Browse** button appears. Click this button to open a

dialog box where you can select the destination blade server.

OK
Applies your selection and closes the dialog box.

Cancel
Closes the dialog box without applying your selection.

5.5.7 Restore Options dialog box (server profiles)

The **Restore Options** dialog box opens when you select **Profiles** in the third step of the **Restore Configuration** wizard. It also opens when you click the **Options...** button in the **Profiles** row. In this dialog box you can specify additional parameters to restore server profiles.

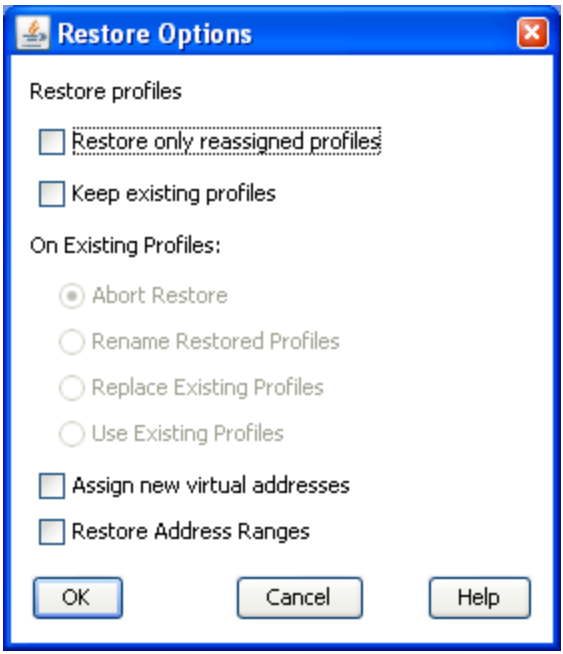

Figure 72: Restore Options dialog box (server profiles)

Restore only reassigned profiles

Specifies that only the profiles that are reassigned to the selected blade servers will be restored (see section "Restoring blade server configurations" on page 306). If you do not select this option, all the profiles saved in the backup file will be restored.

Keep existing profiles

Specifies that the existing profiles remain the same. This option is selected automatically if you select the Restore only reassigned profiles option.

If you do not select this option, all existing profiles are deleted before the configuration is restored.

If you select the **Keep existing profiles** option, select what is to happen with the existing profiles of the same name in **On Existing Profiles**.

Abort Restore

The restore operation is canceled and an error message displayed.

Rename Restored Profiles

The existing profiles are renamed by adding the backup date and possibly also a number to the file names.

Replace Existing Profiles

The existing profiles are replaced by the profiles contained in the backup.

Use Existing Profiles

The profiles are not restored and the existing profiles continue to be used instead.

Assign new virtual addresses

The restored profiles are assigned new virtual addresses.

Restore Address Ranges

Specifies that the origin address ranges will be restored.

This will be necessary if address ranges have been changed since the backup configuration has been saved or if it is created on another management station. This option cannot be selected together with **Keep existing profiles**.

OK
Applies your selection and closes the dialog box.

Cancel
Closes the dialog box without applying your selection.

5.5.8 Select Profile dialog box

In the **Select Profile** dialog box, you can select the required server profile in the tree structure. The area on the right displays information on the selected profile.

To open this dialog box, click **Assign Profile** or select **Assign Profile** in the context menu of the required slot or PRIMERGY rack server on the **Server Configuration** tab.

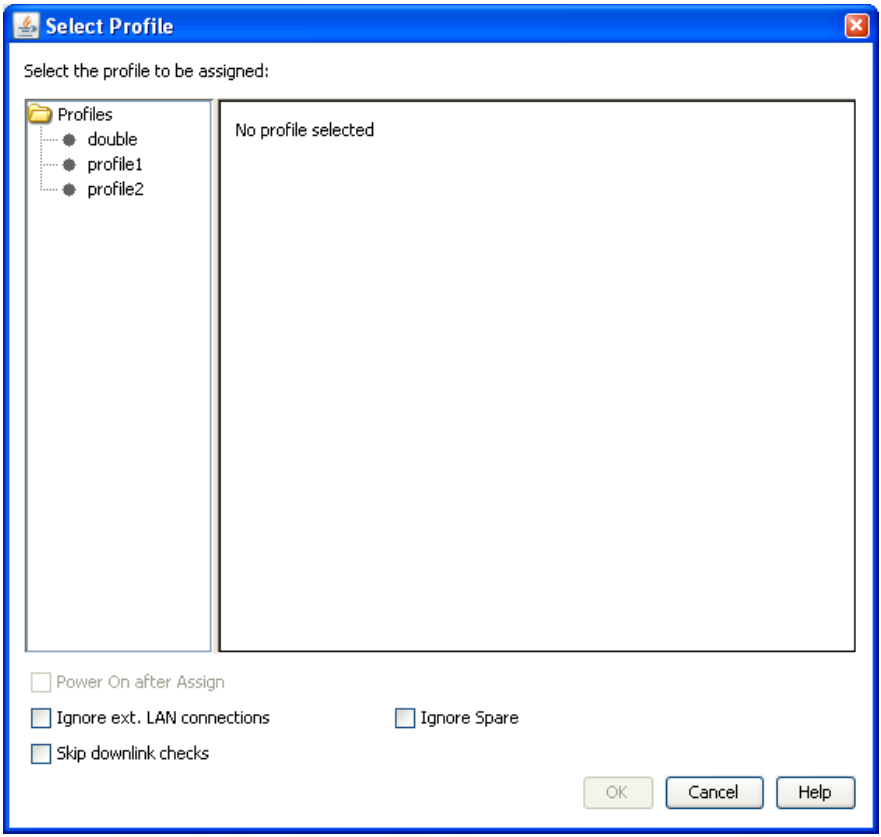

Figure 73: Select Profile dialog box

Power On after Assign

The server is started once the server profile has been assigned.

To prevent further questions when warnings occur, you can state here that you want to assign the server profile even if a warning is issued:

Ignore ext. LAN connections

If you try to assign a server profile with configured networks to a normal LAN switch, you will get an error. With this option you ignore all networks of the server profile. Then no paths are configured and only the virtual addresses are used.

Ignore Spare
No warning during assignment to a spare slot

Skip downlink checks
No warning, even though ports are configured in the server profile for which there are no downlinks in the IBP. If a network is defined for such a LAN port, the server profile cannot be assigned in any case. With this option you can assign a server profile even though no switch is set

OK
The selected server profile is assigned to the slot.

If you have not selected **Ignore ext. LAN connections**, **Ignore Spare**, or **Skip downlink checks**, a corresponding warning is issued in another window. In this case you must confirm that you still want to assign the server profile to this slot.

If the server profile is already assigned to another slot, a corresponding message appears in another window asking whether you wish to continue with the operation. If you confirm this, the previous assignment is deleted and the profile is assigned to the new slot.

If you try to assign a PRIMERGY rack server profile with PCI cards to a blade server slot or a blade server profile to a PRIMERGY rack server, a warning is shown. If you confirm that you really want to assign the profile, only the mezzanine/PCI cards in slot 1 and 2 are regarded (see "VIOM server profile mapping" on page 132).

5.6 Context menus

This section contains a list of the various context menus as well as a brief description of the individual menu items.

5.6.1 Context menus on the Ext. LAN Connections tab

The **Ext. LAN Connections** tab provides a number of context menus.

The context menu of an uplink set contains the following menu items:

New Uplink Set
Define a new uplink set

Edit Uplink Set
Edit an uplink set

Delete
Delete network(s) or uplink set(s)

The context menu of an uplink port can contain the following menu items depending on how it is configured:

Add
Assign port to an uplink set

Add as Backup
Assign port to an uplink set as a backup port

Remove
Remove port assignment from an uplink set

Active
Change configuration of a backup port in an active port

Backup
Change configuration of an active port in a backup port

A network must be selected in the table in order to make the context menu available for an uplink port.

5.6.2 Context menus in the Server Profiles view

Depending on what you select in the **Server Profiles** view of the **ServerView Virtual-IO Manager** window, context menus with different menu items are displayed.

The context menu of the **Profiles** group in the left area of the window contains the following menu item:

New Profile
Create a server profile

The context menu of a profile contains the following menu items:

Edit Profile
Edit selected server profile

Show Profile Details
View definition of a selected server profile

Delete Profile
Delete selected server profile

Copy Profile
Create a copy of the selected server profile

5.6.3 Context menu on the Server Configuration tab

The context menu of the **Server Configuration** tab contains the following menu items:

Assign Profile
Assign server profile to a slot or PRIMERGY rack server

Unassign Profile
Remove assignment of a server profile from a slot or PRIMERGY rack server

Create profile
Create a server profile

Show Profile Details
View definition of the assigned server profile

Update State
Update the power state display, the boot mode and the virtualization status of a server

Inventory Boot

Re-create the inventory table of the server.

During inventory boot the system BIOS assembles the inventory information of the server blade or PRIMERGY rack server hardware as needed by the Virtual-IO Manager and sends it to the management blade or to iRMC, where it is stored.

In some cases a manual execution of the inventory boot is necessary in order to support the new functionality:

- For support of new functionality a new version of the inventory table might be needed. Example: Converged network adapter card (CNA)
- In some cases the new version of the inventory table might not be created automatically for the server. Example: The new version of the inventory table is only supported by a new version of the system BIOS.

You must execute the inventory boot manually, after applying a new version of the system BIOS and iRMC firmware or a new version of the firmware of the optional hardware.

Boot

Start the server(s)

Shutdown

Switch off the server

Failover

Assign the server profile of a slot to a suitable spare slot in the event of a failure.

Video Redirection

Open a new window where the console output of the server is shown.

5.7 General buttons

This section describes the general buttons that you will come across in the Virtual-IO Manager.

5.7.1 Buttons in the area on the left

Server List button

Click **Server List** in the left area of the Virtual-IO Manager to switch to the file tree view according to the ServerView server list.

Profiles button

Click **Profiles** in the left area of the Virtual-IO Manager to switch to the profile view of the defined server profiles.

5.7.2 Button in the area on the right

Click this button to refresh the display on the tab.

5.7.3 General buttons in other dialog boxes

Back button

Click **Back** to return to the previous step of the relevant wizard.

Cancel button

Click **Cancel** to close a wizard/dialog box without saving your changes.

Finish button

Click **Finish** to confirm your entries and exit the relevant wizard.

This button is only active if all the required entries have been made.

Help button

Click **Help** to launch the context-sensitive online help.

Next button

Click **Next** to go to the next step of the relevant wizard.

OK button

Click **OK** to confirm your entries. The dialog box closes.

5.8 Icons

This section contains a list of the VIOM-specific icons and their meaning. These icons are displayed on the **Setup** tab when you click a module in the display.

Icon	Significance
	The module cannot be managed.
	The module has minor configuration problems.
	The module has major configuration problems.

Table 4: Icons on the **Setup** tab

6 Using the Virtual-IO Manager

6.1 Starting the Virtual-IO Manager

You can start the Virtual-IO Manager from the main window of the ServerView Operations Manager:

1. Start the ServerView Operations Manager. For information on starting the ServerView Operations Manager, see the ServerView Operations Manager user guide.
2. Start the Virtual-IO Manager on the start page of the Operations Manager by choosing **Administration** – **Virtual-IO Manager**.

On Windows, you can also launch VIOM from the Windows Start Menu.

1. Choose **Start** – **[ALL] Programs** – **Fujitsu** – **ServerView Suite** – **Virtual-IO Manager** – **Virtual-IO Manager**.

In both cases, the login page of the Fujitsu ServerView Suite Central Authentication Service is displayed. Here, you enter a valid user name and password of a user with **AccessVIOM** privilege. By default, this is a user with **Administrator** role.

The start page of the Virtual-IO Manager is then displayed.

6.2 Closing Virtual-IO Manager

You can close the Virtual-IO Manager by closing the main window.

1. To close the Virtual-IO Manager main window, click on the **Close** icon in the browser window.

6.3 Logging the actions using VIOM

6.3.1 Logging the actions on Windows

The Virtual-IO Manager logs all actions in the Windows Event Log, including logins to the Virtual-IO Manager and all changes to the specific blade server configuration on the Virtual-IO Manager. To do this, the Virtual-IO Manager creates a separate Event Log **ServerView VIOM**, in which it saves its events (see "Event logging" on page 359).

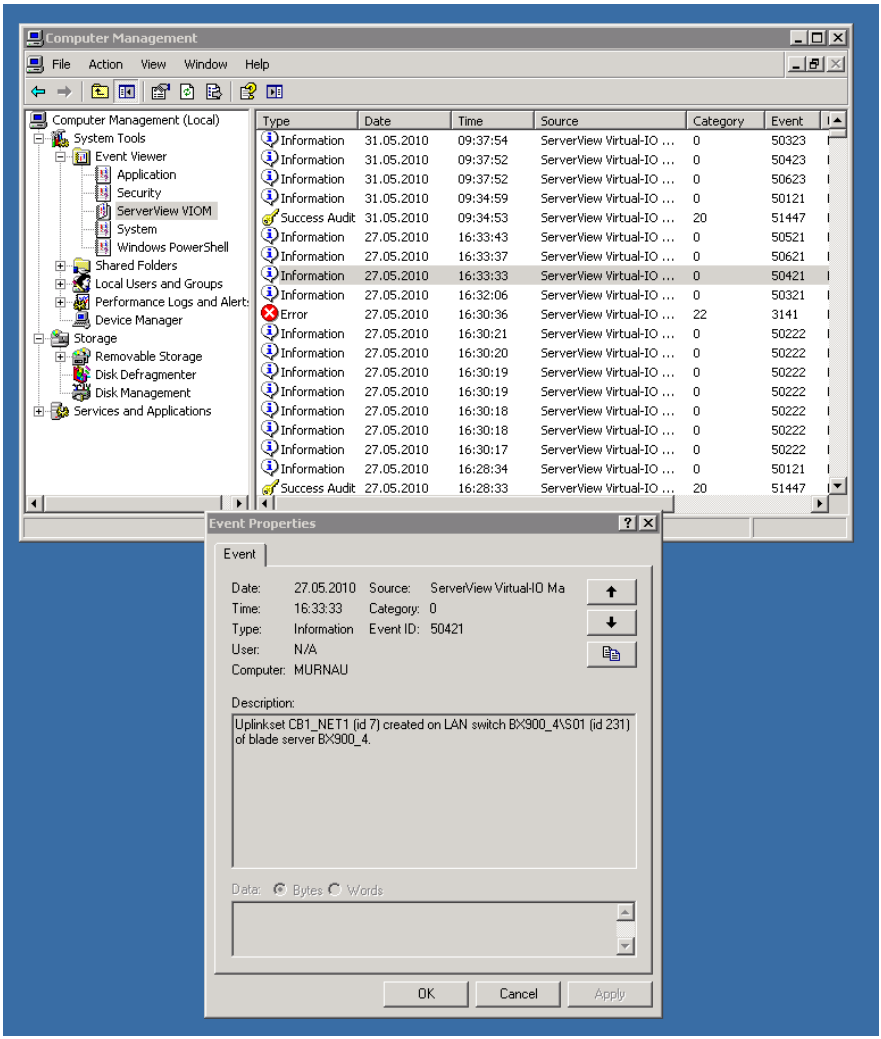

Figure 74: VIOM-specific log file

6.3.2 Logging the actions on Linux

On Linux the actions of VIOM are logged using the Syslog mechanism. All informational messages are logged using the Syslog facility **user**. For example on SLES 11 you will find the VIOM messages in the file **/var/log/user_messages**.

7 Managing servers with VIOM

When you install a new server that can be managed using VIOM and add it to the ServerView server list, then this server is automatically added to the VIOM-specific server group **VIOM Manageable**.

In order to manage the server with VIOM, you first have to add it to the group of servers managed with VIOM (**VIOM Managed**). When moving a server from one group to another, you have to define the access rights as well as the protocols and ports used for the individual modules.

The following sections describe how to activate or deactivate VIOM management of a server and how to set or change the access rights and ports.

Please note: A server may always only be managed by one management station with the Virtual-IO Manager. If, nevertheless, you try to manage a server from a different management station, you receive the message that this server is already being managed by a management station.

You can ignore this warning if you want to manage the VIOM from another management station. In this case, however, you need to ensure that the other installation of the Virtual-IO Manager no longer exists. Otherwise you will have undefined statuses in the managed server.

After the Virtual-IO Manager is reconstructed on the CMS due to a hardware failure, uninstalling and reinstalling, etc., you will need to power off all target servers that are managed by the Virtual-IO Manager, because the networks have to be reconfigured and the server profiles are reassigned.

7.1 Activating management with VIOM

To manage a server with VIOM, you have to select the server from the **VIOM Manageable** server group and add it to the **VIOM Managed** server group. You do this using the **Manage** button on the **Setup** tab.

To manage a blade server, you have to access the management blade of the BX600/BX400/BX900 chassis and the IBP modules in the chassis. To access the management module (MMB) and the IBP modules with VIOM, you have to specify data for authentication during activation, e. g. passwords and possibly also protocols and port numbers.

To manage a PRIMERGY rack server, you have to access the iRMC of this server. To access the iRMC you must specify user name and password.

1. Select the required server. To do this, click the required server in the **VIOM Manageable** server group on the left of the VIOM window.

   The server is displayed on the **Setup** tab. The tab view corresponds to the model and configuration of the selected server.

   

   As long as the server does not belong to the **VIOM Managed** server group, only the **Setup** and **Virtual-IO Manager** tabs are activated.

2. Click the **Manage** button.

   If you are going to manage a blade server, you receive a warning that the previous configuration will be deleted. This means that the default settings following the initial operation of the blade server or even the user-defined settings in the IBP modules will be deleted. For more information on this, see section "VIOM internal operations on blade servers" on page 261.

3. If you click **Yes**, the **Authentication** dialog box opens. This dialog box varies depending on the selection in the table on the **Setup** tab. See "Authentication dialog box (single blade server)" on page 231, "Authentication dialog box (PRIMERGY rack server)" on page 234 and "Authentication dialog box (PRIMERGY rack server and blade server)" on page 236.

   In this dialog box, you enter the user names/passwords for the management blade, for the IBP modules, or for the PRIMERGY rack servers, which VIOM can use to access these components.

   By default, the user name **admin** and password **admin** are set up for all components. In this case, you must only enter the user name and

password for the management blade or the first PRIMERGY rack server and click **Use for all components**. These settings are then applied to all I/O connection blades or all PRIMERGY rack servers.

You should change the default user name and password combination for security reasons. You can have the same user name and password for all components or different ones for each. You must specify this accordingly in this dialog box.

4. Click the **Configure protocols** button if you do not use the default protocols or ports. The **Authentication** dialog box expands so that you can also configure the protocol and ports used by VIOM to access the components.
   1. Select the protocol to be used for communication with the management blade, the I/O connection blades, or PRIMERGY rack servers (**Protocol**).
   2. Specifying the port numbers for the components if **Use default port** is not checked.

   For a standard configuration, these values should not be changed.
5. When you click **OK**, the configuration is applied and the **Authentication** dialog box closes. The tabs in the area on the right are then activated.
6. When you click **Cancel**, you close the **Authentication** dialog box without applying your changes and end the activation process.

7.2 Changing access rights and ports

You can also change the access rights and ports for the components (MMB, IBP, or PRIMERGY rack servers) retrospectively.

1. Select the required server. To do this, click the required server on the left of the VIOM window in the **VIOM Managed** server group.

2. Click **Authentication** on the **Setup** tab. The **Authentication** dialog box opens. See "Authentication dialog box (single blade server)" on page 231, "Authentication dialog box (PRIMERGY rack server)" on page 234.
3. Enter the user names and passwords for the individual components (for the management blade and for all I/O connection blades or for the PRIMERGY rack servers), which VIOM can use to access these components.

   By default, the user name **admin** and password **admin** are set up for all components. In this case, you must only enter the user name and password for the management blade or the first PRIMERGY rack server and click **Use for all components**. These settings are then applied to all I/O connection blades or all PRIMERGY rack servers.

   You should change the default user name and password combination for security reasons. You can have the same user name and password for all components or different ones for each. You must specify this accordingly in this dialog box.
4. Click the **Configure protocols** button, if you do not use the default protocols or ports. The **Authentication** dialog box expands so that you can also configure the protocol and ports used by VIOM to access the components.
   1. Select the protocol to be used for communication with the management blade or the I/O connection blades (**Protocol**).
   2. Specifying the port numbers for the modules if **Use default port** is not checked.

   

   For a standard configuration, these values should not be changed.
5. When you click **OK**, the configuration is applied and the **Authentication** dialog box closes.
6. When you click **Cancel**, the **Authentication** dialog box closes without applying your changes.

7.3 Deactivating management with VIOM

You deactivate the management of servers by moving the servers from the **VIOM Managed** server group to the **VIOM Manageable** server group.

1. Select the required servers. To do this, click the required server on the left of the VIOM window in the **VIOM Managed** server group.
2. Click **Unmanage** on the **Setup** tab.
3. If you are unmanaging a blade server, you will receive a warning that the previous configuration will be deleted if you move the server. The warning also informs you that, before deactivating management with VIOM, you should remove all the LAN cables except the one connected to the first uplink port.

   If you click **Yes**, the previous configuration with VIOM is deleted and the blade server is added to the **VIOM Manageable** server group.

   If you click **No**, you close the window without moving the server.

   For PRIMERGY rack servers: If you are only unmanaging PRIMERGY rack servers, a simple confirmation message is shown.

For more information on the internal operations for deactivating management with VIOM, see section "VIOM internal operations on blade servers" on page 261

7.4 VIOM internal operations on blade servers

This section contains a list of internal operations that are executed when you activate or deactivate management with VIOM for a blade server.

VIOM internal operations during activation

When you activate management with the Virtual-IO Manager for a blade server chassis, the following actions are executed internally:

1. The chassis is indicated as "managed". This happens, on the one hand, in the VIOM database. However, information is also defined in the management blade of the chassis that this chassis is managed by this installation of the Virtual-IO Manager.

   This information in the management blade is always checked before a chassis is added in the VIOM management. If VIOM can recognize from this that a chassis is already managed by another Virtual-IO Manager, you receive a corresponding warning. If, however, you are sure that a chassis is no longer managed by another VIOM installation, you can still add the chassis to the management despite the warning.

2. All IBP modules in this chassis are set to an initial state that has the following properties:

   - The downlink ports of an IBP module no longer have a connection. This means that the server blades are not connected to each other or to an external network.
   - Except for the first uplink port, no uplink port, in other words, none of the other external ports, has an external connection.

     In the case of the LAN connection blades of the BX400/BX900 (currently PY CB Eth Switch/IBP 1Gb 36/8+2, PY CB Eth Switch/IBP 1Gb 36/12, PY CB Eth Switch/IBP 1Gb 18/6, and PY CB Eth Switch/IBP 10 Gb 18/8), which support an IBP mode, no external port (not even the first uplink port) has an external connection.

     The figures below show the standard configuration of the IBP module before activating and the initial configuration after activating VIOM management (**Manage** button).

Standard configuration before activating VIOM management

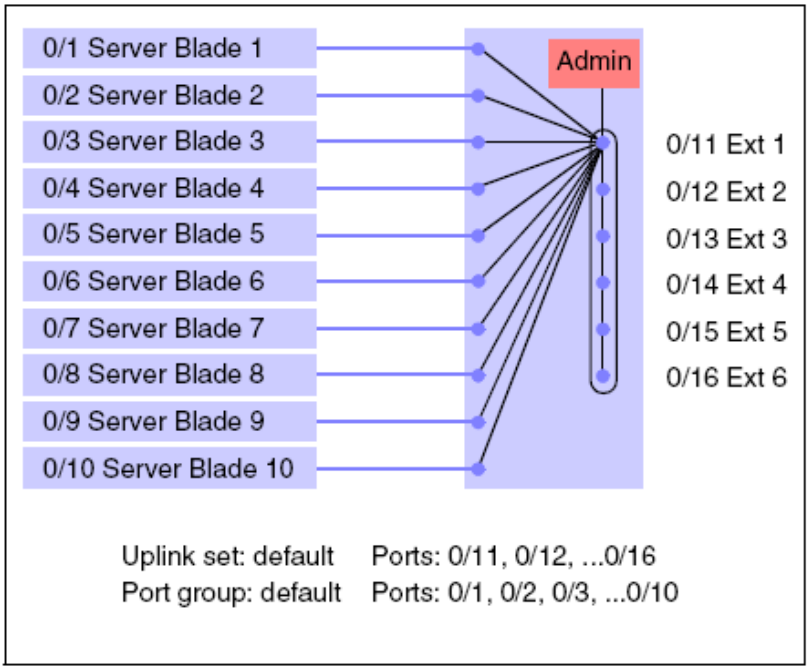

Figure 75: Standard configuration of the IBP 10/6 before activation

In this configuration, all server blades are linked with the standard uplink set. The six uplink ports together form a LAG.

Initial configuration after activating VIOM management

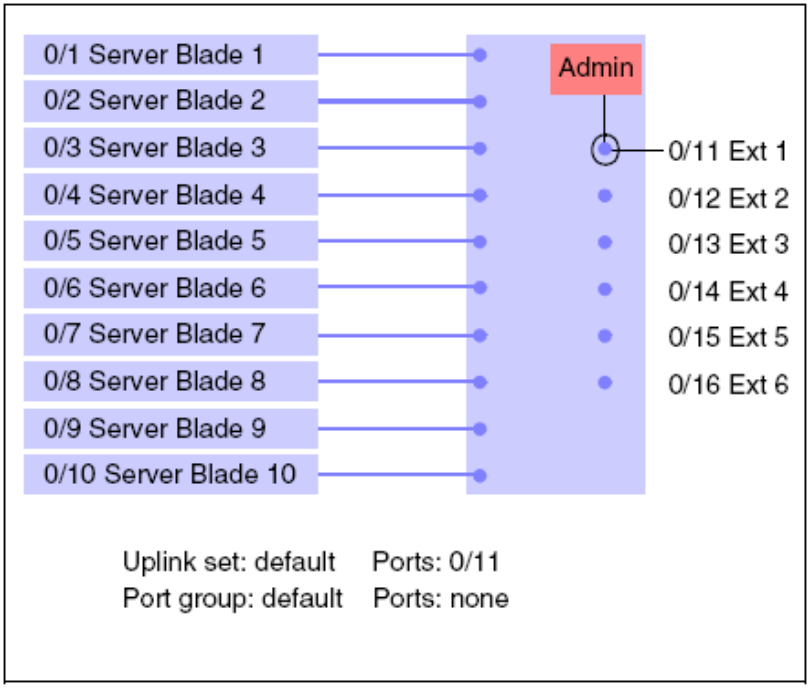

Figure 76: Initial configuration of the IBP 10/6 after activation

After activating the VIOM management (**Manage** button), the server blades are not connected to each other and are not connected to any external networks. Only the first uplink port has an external connection. To enable you to configure the IBP module, the network to which the central management station is connected must be connected to this first uplink port.

In the case of the LAN connection blades of the BX400/BX900, which support an IBP mode (currently PY CB Eth Switch/IBP 1Gb 36/8+2, PY CB Eth Switch/IBP 1Gb 36/12, PY CB Eth Switch/IBP 1Gb 18/6, and PY CB Eth Switch/IBP 10 Gb 18/8), none of the uplink ports has a connection after VIOM management is activated. These modules must be

configured so that they can be accessed from the central management station via their management port.

3. In the management blade, any existing virtualization data (server profile data) is deleted for each server blade slot.

VIOM internal operations during deactivation

When you deactivate management with the Virtual-IO Manager for a blade server chassis, the following actions are executed internally:

1. The initial state (see figure "Standard configuration of the IBP 10/6 before activation" on page 263) is restored for the IBP modules. In detail, this means:
   - First, all defined networks and uplink sets, apart from the standard uplink set, are deleted.
   - The downlink ports are connected to the standard uplink set.
   - All uplink ports are added to the standard uplink set as active ports. In doing this, you have to note that the uplink ports form a LAG. Before you deactivate management with the Virtual-IO Manager, it is essential that you therefore first remove all the LAN cables except for the one connected to the first uplink port.
2. In the management blade, any existing virtualization data (server profile data) is deleted for each server blade slot.
3. The blade server chassis is indicated as "unmanaged". This happens, on the one hand, in the VIOM database. In addition, in the management blade of the blade server chassis, the information that this blade server chassis is managed by this installation of the Virtual-IO Manager is deleted.

7.5 VIOM-internal operations on a PRIMERGY rack server

This section explains the internal operations perfomed when you activate or deactivate management with VIOM for a PRIMERGY rack server.

VIOM-internal operations during activation

When you activate management with the Virtual-IO Manager for a PRIMERGY rack server, the following actions occur internally:

1. The PRIMERGY rack server is indicated as “managed” by this VIOM installation. This is done by assigning a corresponding status to this server node in the VIOM database and also by storing information about the VIOM installation in a VIOM-specific property in the iRMC. This information uniquely identifies the current VIOM installation. (Note: This information is not changed during an update installation of VIOM, but it does change when uninstalling and during a new installation.)

   

   This information in the iRMC is always checked before a server is managed. If VIOM can recognize from this that a PRIMERGY rack server is already managed by another Virtual-IO Manager, you receive a corresponding warning. If, however, you are sure that a server is no longer managed by another VIOM installation, you can still add the server to the management despite the warning.

   VIOM also changes a special VIOM-enabled flag in the iRMC from “not enabled” to “enabled” and a trap destination is stored in the iRMC. The trap destination is one of the required details that VIOM needs to manage a PRIMERGY rack server. The trap destination must be the IP address or server name of the ServerView management station where this Virtual-IO Manager and ServerView Operations Manager are installed. When you specify a server name, the iRMC must be able to resolve this server name into a valid IP address via DNS so that SNMP traps sent by iRMC to the IP address reach the VIOM management station.

The iRMC has limited storage to store multiple trap destinations. This configuration can be modified by other programs or by using the iRMC user interface. When removing the trap destination of the VIOM management station from the iRMC configuration, this might result in incorrect behavior of the server, in the case of a power loss event for this server.

VIOM also needs a username and password in order to authenticate on the iRMC when communicating with the iRMC via RMC protocol (IPMI over LAN).

2. The Virtual-IO Manager performs a special so-called “inventory boot”, which creates an inventory table for this server in the iRMC. This inventory table is used to check the availability of the required I/O devices when assigning a VIOM server profile to this PRIMERGY rack server.

   Therefore you will see that a PRIMERGY rack server is powered on when a server is managed by VIOM. This inventory boot does not boot an operating system; it powers down the server when the inventory table is created and stored in the iRMC.

After these actions the PRIMERGY rack server is integrated into the VIOM management and is ready to be assigned a VIOM server profile. When assigning a VIOM server profile the following actions are performed:

1. VIOM checks that the server is powered off by reading the power state reported by the iRMC.
2. VIOM reads the inventory data stored in the iRMC and checks whether the server supplies the I/O devices configured in the VIOM server profile that is to be assigned to this server.
3. VIOM writes virtualization data to the iRMC. This data contains the virtual addresses for all the configured I/O devices in the VIOM server profile and also the boot information. The virtualization data also contains the port-disable configuration for any I/O device found in the inventory data of this server, for which the server profile does not contain a configuration. VIOM cannot guarantee that all I/O devices that are not configured in the profile are disabled. This only works on devices that

support this functionality.

4. VIOM performs a special so-called init boot that initializes all the I/O devices with the virtualization configuration. This init boot can be avoided and a normal boot performed instead. This option is available in the VIOM user interface when assigning a profile: option **Power On after Assign**. All the VIOM-specific initialization is also performed during a normal boot. The special init boot does not try to boot an operating system, it always powers off a server after the initialization is done.

   When the init boot is successfully completed, the server reports the status **VIOM_SETTING_OK**. This status information is available in the VIOM user interface on the **Server Configuration** tab in the table column **state**. In some situations you may need to manually update this information in order to get the current state. To do this, click the **Update States** button.

   The status **VIOM_SETTING_FAILED** indicates a failure. In this case the server will not be able to boot and a request to power on will result in a power-off status when the VIOM initialization fails. The system event log accessible via the iRMC interface will contain further information about this event.

   The status **POWER_ON_PENDING** indicates that an initial or normal boot is still outstanding and the virtualization configuration is not in effect.

When un-assigning a VIOM server profile the following actions are performed:

1. VIOM checks that the server is powered off by reading the power state reported by the iRMC.
2. VIOM clears the VIOM-specific virtualization data stored in the iRMC. Internally the VIOM server profile is no longer assigned to this PRIMERGY rack server.
3. VIOM performs the special init boot (see also actions when assigning a profile). This is done in order to remove any virtual addresses from the I/O devices so that the original manufacturer addresses (MAC

addresses or WWN addresses) are used from now on. Also the boot device priority setting as defined in the VIOM server profile is no longer valid. The performed init boot results in a power-off state and the server reports the status “Boot without VIOM” which is also displayed in the VIOM user interface on the **Server Configuration** tab in the table column **state**. In some situations you may need to manually update this information in order to get the current state. To do this, click the **Update States** button.

Behavior after AC failure

The VIOM-specific virtualization data and inventory data are stored in the iRMC. This data only remains valid as long as the iRMC is powered on. During a power loss (disconnection of all power cables or power loss in the data center) the iRMC will forget the VIOM virtualization data as well as the VIOM-specific inventory data. When AC power returns and the iRMC has booted again, it still has the following information:

1. This PRIMERGY rack server is managed by VIOM.
2. The VIOM-enabled flag is set.
3. The unique identification of the VIOM installation.
4. The SNMP trap destination as specified by VIOM when managing the server.
5. Whether VIOM virtualization data was stored in the iRMC before power loss or not.

If VIOM virtualization data was stored in the iRMC before power loss and because the iRMC knows that the server was managed by VIOM and that virtualization data was assigned, a special communication process with the VIOM management station will start when the iRMC is booted again after power has returned. This communication is based on sending SNMP traps to the VIOM management station and setting specific status information in the iRMC. For this process to work, the VIOM management station must be running and must be accessible to the iRMC.

As long as the iRMC did not receive virtualization data from the VIOM management station, it will periodically send SNMP traps to the management

station indicating that it is still waiting for these data. This process can be interrupted by disabling the VIOM managed status in iRMC user interface or BIOS interface.

At the end of this process the VIOM inventory data is freshly recreated (inventory boot), the VIOM virtualization data is rewritten by VIOM to the iRMC (also evaluating the new inventory data) and the virtualization data is made active via an init boot. As long as the virtualization data is not available to the server, this server will not be able to boot.

⚠ **This illustrates that when managing PRIMERGY rack servers with VIOM it is very important that the VIOM management station is always running.**

It is advisable to make regular backups of the ServerView Operations Manager database and the VIOM database including transaction logs. These database backups and transaction logs should be stored on high-availability storage media. Regular backups of the OS image of the ServerView management station should also be created so that the VIOM management station can always be recreated in the event of a server crash.

If the server needs to boot after AC failure even though the VIOM management station is not currently available, the VIOM-managed status of this server must be deactivated (see also below). But if this is done, the I/O address virtualization of the server gets lost. The Virtual-IO Manager later detects that this server is no longer VIOM-managed by the VIOM-enabled flag in the iRMC. Because of this, the VIOM-managed status in the Virtual-IO Manager is changed to unmanaged in this case and, if a profile was assigned to this server, it is internally unassigned.

If a VIOM-managed PRIMERGY rack server is moved to a different location that is not controlled by the VIOM management station that originally controlled it, this PRIMERGY rack server cannot boot because it waits endlessly for the virtualization data.

It is best to unmanage the PRIMERGY rack server in VIOM before moving it to the new location. If this is no longer possible, the VIOM-managed status can be deactivated in the iRMC user interface and the system BIOS.

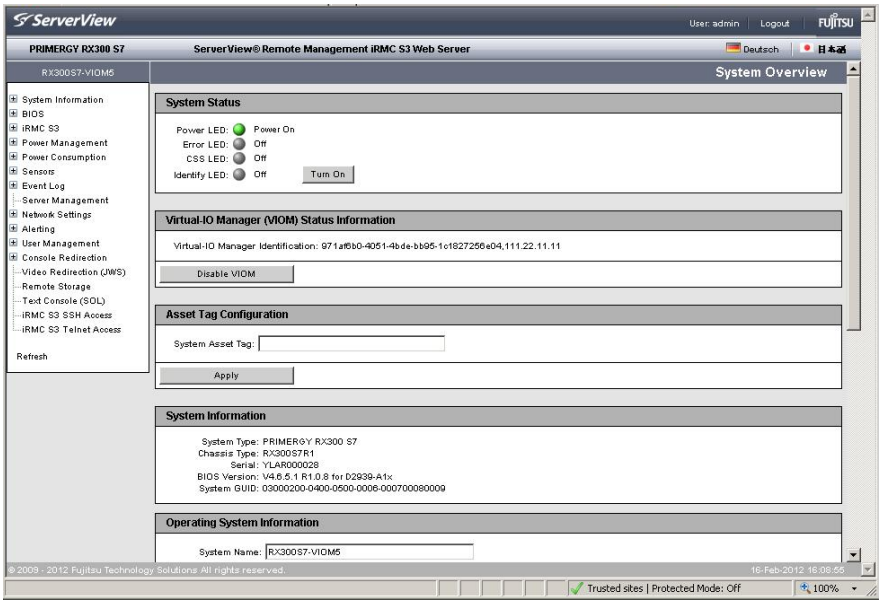

To deactivate the server's VIOM-managed status in the iRMC user interface, click the **Disable VIOM** button in the iRMC view **System Overview**.

To deactivate the server's VIOM-managed status via the system BIOS, open the **Advanced** menu:

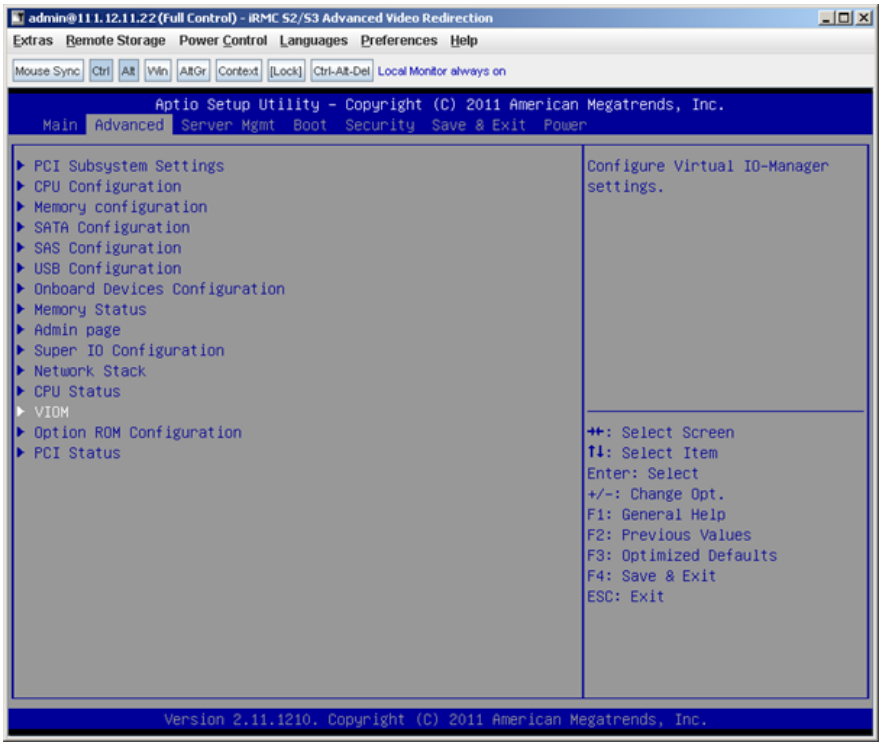

Select **VIOM** in the **Advanced** menu and set the **VIOM-flag** from **Enabled** to **Disabled**:

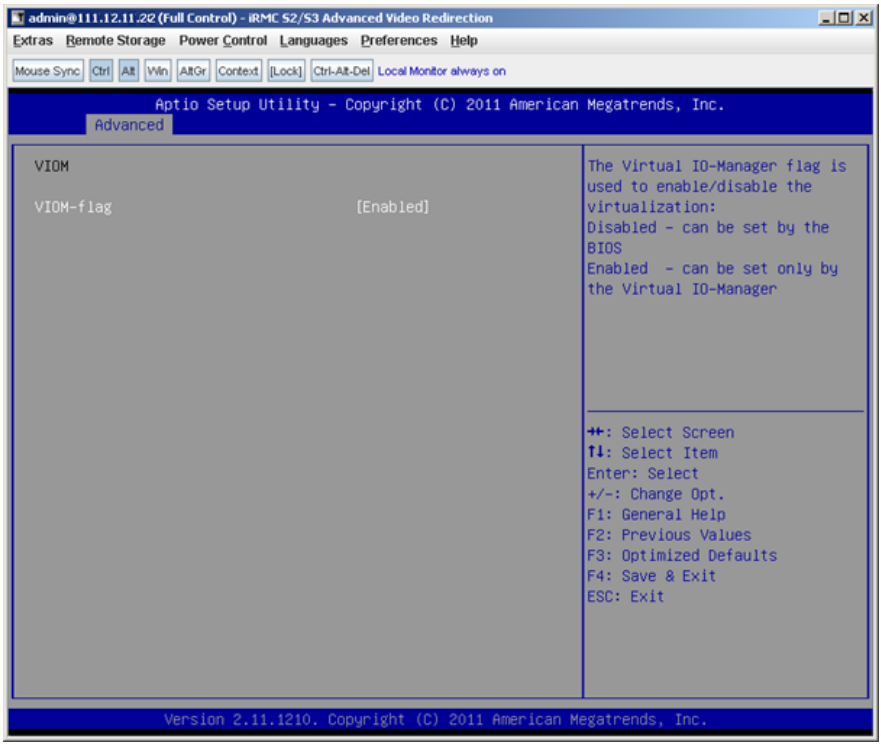

Regardless of whether you deactivate the VIOM-managed status via the iRMC interface or via the BIOS interface, the VIOM-enabled flag is reset, and the information that identifies the VIOM management station, the VIOM virtualization data, and the inventory data are cleared. But please note that the iRMC configuration of the VIOM management station as an SNMP trap destination is not removed by this process. You will need to manually check the list of SNMP trap destinations and delete all unrequired entries. To do this, open the view **Alerting - SNMP Traps** in the iRMC user interface:

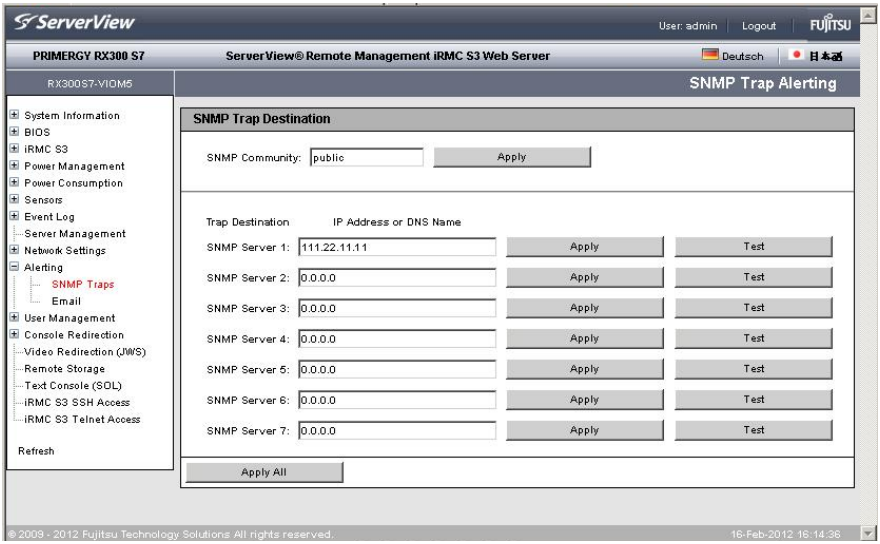

VIOM-internal operations during deactivation

When you deactivate management with the Virtual-IO Manager for a PRIMERGY rack server, the following actions occur internally:

1. VIOM checks that the server is powered off by reading the power state reported by the iRMC.
2. If a VIOM server profile is assigned to this server, the server profile is internally unassigned and the VIOM virtualization data is deleted in the iRMC.
3. The iRMC property that holds the information about the VIOM management station is cleared.
4. The VIOM-enabled flag in the iRMC is reset, indicating that the server is no longer VIOM-managed. This action invalidates the VIOM inventory data in the iRMC.
5. If a VIOM server profile was assigned to this server, an init boot is initiated. With this action the virtualization of I/O addresses is deactivated, so the original manufacturer addresses are activated again. Also the boot priorities set in the server profile are no longer valid.

7.6 Displaying license information

You use the **Show Licenses** button on the **Virtual-IO Manager** tab to request information on the licenses assigned.

1. To do this, click the **Show Licenses** button on the **Virtual-IO Manager** tab. The **Licenses Information** dialog box opens.
2. Click **OK** to close the **Licenses Information** dialog box.

The information displayed includes the following:

- The total number of server profiles that can be assigned using VIOM with the license entered
- The number of server profiles already assigned with the license
- The number of server profiles that may still be assigned with the license

If the permitted total number of licenses has been reached, you cannot assign any more server profiles. Further assignment (**Assign Profile**) is denied and a corresponding error message is displayed. The only two possibilities for still assigning another server profile are

- to unassign another server profile or
- to specify another license (see section "License management" on page 90)

8 Defining network paths (LAN)

This chapter describes how to define network paths on an IBP module (LAN connection blade) using VIOM.

In the Virtual-IO Manager, defining a network path refers to specifying which external ports are used to connect the relevant blade server chassis to which external networks.

Defining these types of network paths on an IBP module includes the following steps:

- Defining an uplink set. An uplink set contains a number of uplink ports. You can combine multiple uplink ports in one uplink set. You can configure the ports as active ports or as backup ports.
- Possibly also defining one or several meaningful network names, which are assigned to the uplink set.

You can configure the IBP modules individually or for a failover configuration, you can generate an identical configuration by copying the definitions of one IBP module to a second module.

You define network paths using the **Ext. LAN Connections** tab. This tab contains two other tabs (**Graphic** and **Details**). On both tabs, the **New**, **Edit**, and **Delete** buttons are available.

- Click **New** to define a new uplink set (see section "Defining an uplink set" on page 278).
- Click **Edit** to edit an existing uplink set (see section "Modifying an uplink set" on page 285).
- Click **Delete** to delete networks or uplink sets (see section "Deleting networks" on page 286).

On the **Details** tab, you can edit existing configurations directly using the table (see section "Modifying an uplink set" on page 285).

Context menus are also available on the tabs for defining network paths. These will be indicated to you at the relevant places.

The **Chassis Configuration** tab provides an overview of the configuration (IBP configuration and server profile assignment) of a blade server (see chapter "Viewing the blade server configuration" on page 301).

8.1 Defining an uplink set

To define an uplink set, follow the steps below:

1. You can start the **Create network** wizard by:
   - clicking **New** on the **Graphic** or **Details** tabs on the **Ext. LAN Connections** tab
   - selecting **New Uplink Set** in the context menu of an uplink set.

   The **Create network** wizard is launched.
2. In the first step of the wizard, you specify what type of uplink set or what network you wish to create. You can choose between:
   - **Internal network**
     Creates an internal network without a connection to an uplink port. This establishes a connection between the server blades (internal IBP connections) without there being a connection to an external network. To find out what other entries you need to make during the further course of the wizard, see section "Defining an internal network" on page 279.
   - **Single network** (selected by default)
     Creates an uplink set with a network. The uplink set and the network have the same name. To find out what other entries you need to make during the further course of the wizard, see section "Defining a single network" on page 280.
   - **VLAN / Service VLAN networks**
     Creates an uplink set to which one network or even multiple networks with a VLAN ID can be assigned. To find out what other entries you need to make during the further course of the wizard, see section "Defining VLAN networks" on page 282.
   - **Dedicated Service network**
     Creates an uplink set with one dedicated service network. To find

out what other entries you need to make during the further course of the wizard, see section "Defining a dedicated service network" on page 285.

You can view the configurations made on the **Graphic** and **Details** tabs.

8.1.1 Defining an internal network

If server blades only need to communicate with each other and, for security reasons, there must be no connection to an external network, then you configure an internal network without a connection to an uplink. In this case, we are dealing with internal IBP connections.

When defining the server profiles and assigning these profiles to the server blade slots, you specify which server blades are to be connected via which LAN ports via an internal network (see chapter "Defining and assigning server profiles" on page 289).

To define an internal network, follow the steps below:

1. Click on the **Ext. LAN Connections** tab in in the work area on the left.
2. Start the **Create network** wizard by:
   - clicking **New** on the **Graphic** or **Details** tabs
   - selecting **New Uplink Set** in the context menu of an uplink set.

   The **Create network** wizard is launched.
3. In the first step of the **Create network** wizard, select **Internal network** and confirm your entry by clicking **Next**. The second step of the **Create network** wizard opens.
4. Enter the required name for the internal network in **Name of network**.
5. Confirm this by clicking **Finish**. The internal network is configured with the specified name.

You cannot edit an internal network. To find out how to delete an internal network, see "Deleting networks" on page 286".

8.1.2 Defining a single network

In VIOM, a single network refers to a network in which a server blade can communicate externally via one or several uplink ports. The uplink ports are grouped in an uplink set. The network is given the same name as the uplink set.

By grouping several active uplink ports in an uplink set and assigning an uplink set to a network, a link aggregation group (LAG) can be formed. Since it is possible to use several connections simultaneously in a LAG, you can achieve a higher transmission speed and greater level of reliability in the network.

A static LAG is formed by default, i. e. the LACP protocol is not supported. Please note that the ports on the external switch, which are connected to the IBP ports that form a LAG, have to form a static LAG if the LAG on the IBP forms a static LAG.

You can optionally enable the LACP mode when creating a network. If the LAG on the IBP is configured to use LACP, the LAG on the external switch must also be configured to use LACP.

Note: The Link Aggregation Control Protocol (LACP) is defined in IEEE 802.3ad, which allows dynamic trunking of two or more network connections between two switches.

For a single network, no VLAN IDs are taken into account, i. e. all packages are allowed through irrespective of the VLAN ID.

To define an single network, follow the steps below:

1. Click on the **Ext. LAN Connections** tab in in the work area on the left.
2. Start the **Create network** wizard by:
   - clicking **New** on the **Graphic** or **Details** tabs
   - selecting **New Uplink Set** in the context menu of an uplink set.

   The **Create network** wizard is launched.

3. In the first step of the **Create network** wizard, select **Single network** and confirm your entry by clicking **Next**. The second step of the **Create network** wizard opens.
4. Enter the required name for the uplink set in **Name of uplink set**. The network automatically has the same name.
5. Under **Ports for the uplink set**, select the uplink ports that should belong to this uplink set. An uplink port can be included in the port group as an active port or as a backup port. An uplink set must contain at least one active port.

   All active uplink ports and all backup ports each form a LAG.

   If a network needs to be configured so that the IBP module switches to another port if any problems arise, then besides configuring the active ports, you also have to configure one or several ports as backup ports. If there is no longer a connection for any of the active ports of an uplink set, the backup ports are activated and the ports that were active until this point are deactivated.

   Assign the required ports of the IBP to the uplink set by opening the context menu of each relevant port.

   

   For a LAG configuration, you have to configure at least two active ports.

   

   With PY CB Eth Switch/IBP 1Gb 36/8+2 connection blades, 1GB uplinks and 10 Gb uplinks must not be mixed. If you try this, the **Finish** button will not be enabled.
6. Select additional options, such as

   Activate Port Backup: Switch to a backup port if an error occurs in the active port

   Linkdown propagation: Send a linkdown event if both the active ports and the backup ports fail. The linkdown event triggers a failover process on the server blade if configured accordingly.

7. If the uplink set is created on a PY CB Eth Switch/IBP 10 Gb 18/8, click **Next** to confirm your entries. The next step of the **Create network** wizard opens. In this step, you can specify the CNA parameters.
8. Click **Finish** to confirm your entries. A network is created with the configured uplink set.

The new network is added to the tables on the **Graphic** and **Details** tabs.

For further information on defining networks, see section "Defining networks (LAN) (for blade servers only)" on page 30.

8.1.3 Defining VLAN networks

When you define VLAN networks, you define an uplink set, referred to as "shared uplink set" in the following, which the various VLAN networks with different VLAN IDs share. As each network within a shared uplink set is given a unique VLAN ID, these networks are completely separate from one another.

By grouping several active uplink ports in an uplink set and assigning an uplink set to a network, a link aggregation group (LAG) can be formed. Since it is possible to use several connections simultaneously in a LAG, you can achieve a higher transmission speed and greater level of reliability in the network.

In this case, a static LAG is formed, i. e., the LACP protocol is not supported. Please note that even the ports on the external switch, which are connected with the IBP ports that form a LAG, also have to form a static LAG.

To define an VLAN networks, follow the steps below:

1. Click on the **Ext. LAN Connections** tab in in the work area on the left.
2. Start the **Create network** wizard by:
   - clicking **New** on the **Graphic** or **Details** tabs
   - selecting **New Uplink Set** in the context menu of an uplink set.

   The **Create network** wizard is launched.

3. In the first step of the **Create network** wizard, select **VLAN / Service VLAN networks** and confirm your entry by clicking **Next**. The second step of the **Create network** wizard opens.
4. Enter the required name for the uplink set in **Name of uplink set**.
5. Under **Ports for the uplink set**, select the uplink ports that should belong to this uplink set. An uplink port can be included in the port group as an active port or as a backup port. An uplink set must contain at least one active port.

   All active uplink ports and all backup ports each form a LAG.

   If a network needs to be configured so that the IBP module switches to another port if any problems arise, then besides configuring the active ports, you also have to configure one or several ports as backup ports. If there is no longer a connection for any of the active ports of an uplink set, the backup ports are activated and the ports that were active until this point are deactivated.

   Assign the required ports of the IBP to the uplink set by opening the context menu of each relevant port.

   

   For a LAG configuration, you have to configure at least two active ports.

   

   With PY CB Eth Switch/IBP 1Gb 36/8+2 connection blades, 1GB uplinks and 10 Gb uplinks must not be mixed. If you try this, the **Next** and **Finish** buttons will not be enabled.
6. Select additional options, such as

   Activate Port Backup: Switch to a backup port if an error occurs in the active port

   Linkdown propagation: Send a linkdown event if both the active ports and the backup ports fail. The linkdown event triggers a failover process on the server blade if configured accordingly.
7. Click **Next** to confirm your entries. The last step of the **Create network** wizard opens. In this step, you assign meaningful names for the net-

works defined via VLAN IDs, which are to be assigned to the uplink set defined in the previous step.

8. Enter the meaningful name of the network to be used by the uplink set as well as the unique VLAN ID of the network within the shared uplink set.

   The networks on an uplink set must have different VLAN IDs. In contrast, however, two networks on two different uplink sets can use the same VLAN ID. These networks are still completely separate from one another.

9. Optional :
   - Specificy a native VLAN network. All packages that do not contain a VLAN ID will be allowed through this connection.
   - Select a network as service VLAN.

10. Click **Add**, the virtual LAN network is added to the table below.
11. Define any other required networks that are to use the port group, and confirm each of these with **Add**.
12. You can also define a network as a native VLAN in the table retrospectively. To do this, select the checkbox in the corresponding row under **Native VLAN** in the table.
13. You can also change the VLAN ID of a network in the table retrospectively. To do this, click in the corresponding row under **VLAN Id** in the table and edit the VLAN ID.
14. You can also define a network as a service VLAN in the table retrospectively. To do this, select the checkbox in the corresponding row under **Service VLAN** in the table.
15. You can also delete VLAN networks from the list. To do this, select one of the VLAN networks in the table and click **Delete.**
16. When you have defined all the networks you need, click **Finish** to confirm your entries. The networks added to the table are created with the configured uplink set.

The new defined VLAN networks are added to the tables on the **Graphic** and **Details** tabs.

For further information on defining networks, see section "Defining networks (LAN) (for blade servers only)" on page 30.

8.1.4 Defining a dedicated service network

Defining a dedicated service network is mostly the same as defining a single network. For further information, see section "Defining a single network" on page 280

1. Click **New** on the **Graphic** or **Details** tabs or select **New Uplink Set** in the context menu to define a new dedicated service network. The **Create network** wizard is launched.
2. In the first step of the **Create network** wizard, select **Dedicated service** network and confirm your entry by clicking **Next**. The second step of the **Create network** wizard opens.
3. The following steps are the same steps as the steps for defining single networks.
4. In addition to the steps for defining a single network, a Service VLAN ID must be specified. Enter the required Service VLAN ID.
5. Once you have made all your entries, confirm these with **Finish**. When you click **Finish**, a dedicated service network is created with the configured uplink set.

The new network is added to the tables on the **Graphic** and **Details** tabs. A dedicated service network contains always the value **service (vlan-id)** in the **VLAN id** column.

8.2 Modifying an uplink set

You can modify a defined uplink set via the **Graphic** and **Details** tabs on the **Ext. LAN Connections** tab.

1. Select the corresponding uplink set in the table on the **Graphic** or **Details** tab.
2. Click **Edit**. The **Edit Uplink Set** wizard is launched in which you can change the configuration.

The name of the uplink set cannot be changed; the field is inactive. The type of the uplink set cannot be changed either.

The **Edit Uplink Set** wizard is the same as the **Create network** wizard (see section "Defining an uplink set" on page 278).

Modifying an uplink set directly in the table

On the **Graphic** and **Details** tabs, you can modify the existing configurations directly via the table (with the exception of the **Chassis** and **IBP** columns on the **Details** tab). The procedure varies according to the column:

- In the **VLAN** column: By double-clicking an entry in the table, a dialog box opens in which you can make your changes.
- In the **Uplink Ports** and **Backup Ports** columns (both only available on the **Details** tab):

  When changing the ports, enter the ports separated by a comma (,). You can define related areas using the minus sign (-), e. g. 11-13 is the same as 11,12,13. When changing the ports in a switch stack, take care to change the value in the appropriate row.
- In the **Linkdown Prop.**, **Port Backup**, **LACP**, and **IGMP** columns on the **Details** tab: Click the check box to change the existing configuration.
- In the port view of the **Graphic** tab: Select the changes in the context menu of the corresponding port.

8.3 Deleting networks

If several networks are assigned to a shared uplink set, you have the option of deleting networks. You can delete networks using either the **Graphic** or **Details** tab.

1. Select the corresponding networks in the table on the **Graphic** or **Details** tab.
2. Click the **Delete** button or select **Delete** in the context menu.

A message appears asking whether you wish to delete the corresponding network.

3. If you confirm this, the networks will be deleted.

If you delete all networks or the last network of an uplink set, the related uplink set is also deleted.

8.4 Copying an IBP configuration

You have the option of copying the defined uplink sets and networks to a second IBP module. This can be useful if you use a teaming configuration for failure safety on the server blades. If the ports of a server blade that are connected to the first IBP each form a LAN team with the corresponding ports connected to the second IBP, and these ports are connected to the same networks, then in the case that the teaming software triggers a failover, the LAN connection remains in the same networks.

If you wish to use the failure safety, then the following requirements must be met:

1. The uplink sets must be defined with the Linkdown propagation function.
2. A LAN team must be configured on the server blade.
3. The network on the second IBP module must be configured accordingly.

The following describes how to create a corresponding IBP configuration by making a copy:

1. Activate the **Graphic** tab on the **Ext. LAN Connections** tab.
2. Select the IBP whose configuration should be copied and click **Copy**.
3. Then select the IBP to which to copy and press the **Paste** button.
4. If you confirm the query, all definitions of the one IBP module are copied to the other IBP module.

- VIOM always copies the entire configuration irrespective of what is selected in the tables.
- If the target IBP has fewer uplink ports than the source IBP, all uplink ports which does not exist on the target are removed from uplink sets.

8.5 Copying configuration

You have the option of copying all defined uplink sets and networks of a chassis to another chassis. The following describes how to create a copy of all networks of a chassis on another chassis:

1. Activate the **Ext. LAN Connections** tab.
2. Click **Copy Configurations**.
3. In the **Select Target Chassis** dialog box, select the chassis to which the configuration should be copied and click **OK**.

   After affirmation to overwrite the existing configuration, copying starts.

The source and the target chassis both should have the same IBP types in the same slots. Otherwise the copy will probably fail.

9 Defining and assigning server profiles

To use address virtualization or to enable a server blade to use a defined network, you must:

1. Define a server profile
2. Assign the profile to a slot or a PRIMERGY rack server

A server profile contains a set of parameters, which contains the VIOM-specific input/output parameters and the input/output connections.

You can define a server profile in two different ways:

- Using the **Server Profiles** view
- Using the **Server Configuration** tab

The server profile is then stored in the server profile repository.

You activate a server profile by assigning it to a slot in the blade server or to a PRIMERGY rack server.

The following functions are available in the profile view:

- You define server profiles by clicking **New** (see section "Defining server profiles" on page 290).
- You modify server profiles by clicking **Edit**(see section "Modifying server profiles" on page 294).
- You delete server profiles by clicking **Delete**(see section "Deleting server profiles" on page 296).

You can also access these functions using various context menus. These will be pointed out to you at the relevant places. You can also copy a server profile using the context menu of a profile (see section "Copying server profiles" on page 295 ").

You assign a server profile to a slot or to a PRIMERGY rack server on the **Server Configuration** tab using the **Assign profile** button or **Assign Profile** in the context menu of the required slot or PRIMERGY rack server (see section "Assigning server profiles" on page 296).

An overview of the defined server profiles is displayed via the profile view on the left of the **ServerView Virtual-IO Manager** window. In this case, a table containing an overview of the profiles defined up to now is displayed on the right of the window.

9.1 Defining server profiles

To define a server profile the **Create Server Profile** wizard is used. This wizard can be opened from two views:

Profiles view

1. Switch to the **Profiles** view in the area on the left of the **ServerView Virtual-IO Manager** window. Click **Profiles**, if applicable.

   If no profiles have been created yet, the **Profiles** group is empty.

2. Click the **New** button in the area on the right or select **New Profile** from the **Profiles** group context menu to launch the wizard for defining a server profile.

ServerList view

1. Switch to the view of the servers managed by VIOM in the area on the left using the **Server List** button.
2. Switch to the **Server Configuration** tab on the right.
3. Select **Create Profile** from the context menu of a table entry.

When the **Create Server Profile** wizard is launched, proceed as follows:

1. Enter the name of the profile in the first step of the wizard (**Name** step). If a profile already exists with this name or the name is invalid, the name is marked in red.
2. Select the type of the profile.
3. Optional: Select a server model under **Preset values to use for server type**.
4. Optional: Provide a comment for the profile.

5. Once you have entered all the required data, click **Next** to go to the **Configure Cards** step of the wizard.
6. Select the type and number of onboard IO channel in the **Configure Cards** step.
7. Optional: Add mezzanine/PCI cards to the appropriate slots and select the type and number of IO channels for each card. You can move a card to another slot by using the arrow buttons on the right.

   The number of onboard ports cannot exceed the maximum possible value for the selected server model. Furthermore, no more mezzanine/PCI cards can be specified than are supported by the selected model.

8. If you have configured the number and type of ports and cards, click **Next** to go to the **IO-Channels** step.
9. Select **Use virtual addresses** to use virtual MAC addresses and WWN addresses with this profile. You can enter these addresses in **Virtual Addresses** step within this wizard or VIOM can assign these addresses automatically.
10. Select **Disable Boot Menu Usage (F12)** to prevent your VIOM boot settings from being overwritten on your local computer.
11. For CNA IO-channels, select the number of physical functions using the  and  buttons and select the type of each physical function using the dropdown-list.
12. **SMUX setting** is only enabled on blade server profiles if a second mezzanine card is defined as a LAN mezzanine card. Here, you can define the fabric to which the card is routed.
13. The upper table displays the onboard LAN ports (up to 6). Another table is displayed for each configured mezzanine or PCI card.
    1. You can specify a network for each LAN or CNA port in a blade server profile under the **Network** column.

If you want to use the profile on blade servers with IBP modules, you can specify a network for each LAN port. If you work with blade servers that have non-VIOM-capable LAN modules (Open Fabric mode), do not specify a network as it is not possible to define networks on these modules.

To enter a network name, click the table cell to switch to edit mode. You can also open a network selection dialog box via the "..." button. In this dialog box you can select a managed blade server chassis from the selection list. The networks defined for this chassis are then displayed. To select a network, double-click the name or select the name and click the **Add** button.

Make sure that the networks entered here are/will be configured before the profile is activated on the corresponding blade server.

As long as a network does not yet exist, the server profile can be created with this network, but cannot yet be assigned to a slot. If you wish to exit the network selection without selecting a network, click another input field.

2. You can specify tagged VLAN networks for each LAN or CNA port in a blade server profile. If you specify more than one tagged VLAN network for a port, the names must be separated by commas. If you use the network selection box, the name of the chosen network is added to the **Tagged VLAN** column if you use the **Add tagged** button.
3. You can specify service networks for each LAN or CNA port in a blade server profile. If you specify more than one service network for a port, the names must be separated by commas. If you use the network selection box, the name of the chosen network is replaced in the **Network** column or added to the **Service** column depending on the type of the selected network.

14. To configure the port as the boot device, select **PXE boot**, **iSCSI boot**, or **SAN boot** from the selection list under **Boot**. If you configure an

iSCSI boot device or SAN boot device, you must specify additional boot settings in the next step.

15. Click **Next** to access another step in this wizard. The next steps depend on what you configured in the **IO-Channels** step.

16. If you configured several boot devices or one boot device that requires additional parameters, the next step is the **Boot Parameter** step.

    In this step, you specify the boot order and boot parameters for each port.

    You specify the boot order by using the arrow buttons on the right to move the boot devices up or down.

    The boot parameters for each port are specified in the corresponding field. You have to specify additional parameters for an iSCSI boot (in the case of LAN ports) and SAN boot (in the case of Fibre Channel ports). You do not have to specify any additional parameters for a PXE boot (in the case of LAN ports).

17. If you configured at least one CNA IO-channel, another step in this wizard is the **CNA Parameter** step. In this step, you specify the CNA parameter for all physical functions.

18. If you selected **Use virtual addresses** in the **IO-Channels** step of this wizard, another step in this wizard is the **Virtual Addresses** step.

    You specify the virtual addresses in this step. For each address, you can enter a virtual address or select **Allocate virtual address**. In the latter case, VIOM automatically assigns a virtual address after you exit the wizard.

    

    Automatic assignment is only possible if you specified address ranges when you installed VIOM (see "Installing the Virtual-IO Manager on a Windows-based CMS" on page 58.

    Next to each virtual address, you see the status of the address.

    If you specify an address, you can select **Ignore range** to disable the check whether the address is in the range that you specified during the installation of VIOM.

19. Once you have entered all the required data, click **Next** to go to the **Confirm** step of the wizard.
20. Use this step to check the entries you have made once again.
21. If the server profile definition is okay, click **Finish** to exit the wizard. In this case, the server profile is created and saved in the server profile repository.

The new server profile is now shown in the **Server Profiles** view.

9.2 Viewing server profiles

You can view a server profile definition. You do this as follows:

1. Switch to the **Server Profiles** view in the area on the left of the **ServerView Virtual-IO Manager** window. Click **Profiles**, if applicable.
2. Select the required profile in the area on the right of the **ServerView Virtual-IO Manager** window. Then click **Details** or select **Show Details** in the context menu of the selected server profile.
3. The **Server profiles <profile name>** dialog box opens to provide information on the selected server profile.

For server profiles that are assigned, you can also open the information window on the **Server Configuration** tab. To do this, select a slot or PRIMERGY rack server that has a server profile assigned to it, and then select **Show Profile Details** in the context menu.

9.3 Modifying server profiles

You can modify server profiles retrospectively.

1. To do this, switch to the **Server Profiles** view in the area on the left of the **ServerView Virtual-IO Manager** window. Click **Profiles**, if applicable.

2. Select the required profile in the area on the right of the **ServerView Virtual-IO Manager** window. Then click **Edit** or select **Edit Profile** from context menu.
3. The **Edit Server Profile** wizard is launched in which you can change the existing server profile definition.

The **Edit Server Profile** wizard is the same as the **Create Server Profile** wizard. For a more detailed description of the wizard, see section "Defining server profiles" on page 290.

You can only modify server profiles that are not assigned. If this is not the case, the **Edit** button or **Edit Profile** menu item is inactive.

Before inserting a server blade with new optional hardware, you should unassign the profile and modify it.

Before detaching the PCI card virtualized from the rack servers, you should unassign the profile.

9.4 Copying server profiles

You can create several similar server profiles by creating copies of an existing server profile.

1. To do this, switch to the **Server Profiles** view in the area on the left of the **ServerView Virtual-IO Manager** window. Click **Profiles**, if applicable.
2. Select the required profile.
3. Select **Copy Profile** from the context menu of the selected server profile. A copy of the relevant server profile is created.

The copy of a server profile is saved in the server profile repository under the same name with the suffix **_1**. If you make several copies, the suffix is incremented. The copy gets a new name and new virtual addresses are assigned. All other properties (e.g. boot parameter) are equal in the original and the copy.

You can only copy server profiles if you specified address ranges for the virtual addressed when you installed VIOM.

9.5 Deleting server profiles

You can delete server profiles you no longer need.

1. To do this, switch to the **Server Profiles** view in the area on the left of the **ServerView Virtual-IO Manager** window. Click **Profiles**, if applicable.
2. Select the required profiles in the area on the right.
3. Click **Delete** or select **Delete Profile** in the context menu of the selected server profile.
4. A message appears asking if you wish to delete the corresponding server profiles.
5. If you confirm this, the server profiles will be deleted.

You can only delete server profiles that are not assigned. If this is not the case, the **Delete** button or **Delete Profile** menu item is inactive.

9.6 Assigning server profiles

You use the **Server Configuration** tab to assign server profiles to individual slots or PRIMERGY rack servers.

The number of possible assigns depends on the registered licenses. Each license contains a count which allows a certain number of assigns. If the assign counts of all licenses are used up, no further assign is possible. Only one assign count is allocated per profile assign even for multislot profiles.

1. Click **Server List**, if applicable, to switch to the server list view in the area on the left.

2. In the tree structure on the left, select the corresponding blade server or PRIMERGY rack server from the **VIOM Managed** group.
3. Switch to the **Server Configuration** tab in the area on the right.
4. Select the required slot of the blade server or the required PRIMERGY rack server in the table.

   For blade servers, you can also assign the server profile to an empty slot.

5. You can only assign a server profile to a server blade or PRIMERGY rack server if the server is switched off (power off). You can see whether the state of the server is correct in the **State** column.

   If the state of the server is not **off**, you can switch it off by clicking **Shutdown** or by selecting **Shutdown** in the context menu:

   1. Click **Shutdown** in the context menu or click the **Shutdown** button.
   2. In the next dialog box, select the type of shutdown (**Graceful Shutdown** or **Forced Power Off**).
   3. If you confirm your selection, the server is switched off.

   You can update the display in the **State** column by clicking **Update States** button or with **Update State** in the context menu.

6. Click **Assign Profile** or select **Assign Profile** in the context menu of the required slot. The **Select Profile** dialog box opens.
7. In this dialog box, select the required server profile in the tree structure. The area on the right displays information on the selected profile.

   To prevent further questions when warnings occur, you can state here that you want to assign the server profile even if a warning is issued:

   - **Ignore ext. LAN connections**
   - **Ignore Spare**
   - **Skip downlink checks**

8. Confirm your selection with **OK**. The selected server profile is assigned to the slot.

   If you have not selected **Ignore ext. LAN connections**, **Ignore Spare** or **Skip downlink checks**, a corresponding warning can be issued in another dialog box. In this case you must confirm that you still want to assign the server profile to this slot.

   If the server profile is already assigned to another slot or PRIMERGY rack server, a corresponding message appears in another dialog box asking whether you wish to continue with the operation. If you confirm this, the previous assignment is deleted and the profile is assigned to the new slot or PRIMERGY rack server.

   If you try to assign a PRIMERGY rack server profile with PCI cards to a blade server slot or a blade server profile to a PRIMERGY rack server, a warnig is shown. If you confirm that you really want to assign the profile, only the mezzanine/PCI cards in slot 1 and 2 are regarded (see "VIOM server profile mapping" on page 132).

9.7 Deleting profile assignments

You use the **Server Configuration** tab to delete the assignment of a server profile from an individual slot or PRIMERGY rack server.

1. Click **Server List**, if applicable, to switch to the server list view in the area on the left.
2. In the tree structure on the left, select the corresponding server from the **VIOM Managed** group.
3. Switch to the **Server Configuration** tab in the area on the right.
4. Select the corresponding slot of the blade server or the PRIMERGY rack server in the table.

5. You can only deactivate a server profile if the corresponding server is switched off (power off). You can see whether the server is switched off in the **State** column.

   If the state of the server is not **off**, you can switch it off by clicking **Shutdown** or by selecting **Shutdown** in the context menu.

   

   You can update the display in the **State** column by clicking **Update States** button or with **Update State** in the context menu.

6. Click **Unassign Profile** or select **Unassign Profile** in the context menu of the required slot.

The assignment is deleted and no server profile is assigned to the corresponding slot or PRIMERGY rack server. The display on the **Server Configuration** tab is updated accordingly.

10 Viewing the blade server configuration

You view the blade server configuration via the **Chassis configuration** tab.

1. Click the **Chassis configuration** tab.The **Chassis configuration** tab contains a schematic display of the existing configuration of the blade server.
2. Click the **Update States** button to update the display of the on/off status of the server blades in the rectangle of the slots.
3. Click an uplink port, a network or a bay to select it. The elements associated with the element you have selected are then highlighted.

11 Saving and restoring

You can save your blade server configuration and server profiles in files and restore them later. Backups like these are useful, for example, if you want to use previous configurations after a new installation.

You can store these backup files both on the management station and locally on the computer on which the Web GUI runs.

These backup files contain the configurations (networks, uplink sets, assigned profiles, spare slot definitions) of one or several chassis and/or profiles.

You save and restore backup files as well as delete them on the management station using a wizard. You launch the wizard using the **Configuration Backup / Restore** button on the **Virtual-IO Manager** tab.

11.1 Saving your configuration and server profiles

The starting point for backing up your configuration and server profiles is the **Virtual-IO Manager** tab:

1. Launch the wizard using the **Configuration Backup / Restore** button.
2. Select **Save Configuration** in the first step of the wizard to save the configuration in a file.
3. Click **Next** to go to the **Select File** step.
4. Select the computer on which you want to save the configuration:
   - Select **Save on Management Station** to save the configuration to the management station.
   - Select **Save local** to save the configuration in a file on the computer on which the GUI runs.
5. When you select **Save on Management Station**, you save the configuration to the central management system. There are two ways of specifying the backup file:

- Select an existing file from the list that will then be overwritten.
- Enter the name for the backup file directly in the input field.You can also write in subdirectories by specifying the entire path name (e. g. **directory/file**). If the required directories do not yet exist, they are created automatically. The backup files are assigned the **.xml** suffix automatically if you have not specified one.

6. Select **Save local** to save the configuration in a file on the computer on which the GUI runs. There are two ways of specifying the backup file:
   - Enter the name for the backup file directly in the input field. On Windows, if you do not specify a complete path including drive letter, the file will be saved on the desktop.
   - Click **Browse** to open the file selection dialog box in which you can navigate to the desired folder and select an existing backup file or specify the name of the backup file. If the file exists, it will be overwritten.
7. Once you have specified a backup file, click **Next** to go to the **Select Data** step of the wizard.
8. In **Select servers to save**, select the servers whose configuration you wish to save. Do not select any servers if you only wish to save profiles.
9. Select **Save Profiles** to specify whether server profiles are to be saved too.

It is essential that you save the profiles along with the configuration if the assigned profiles are to be reassigned once the configuration has been restored.

11.2 Restoring the configuration

The starting point for restoring a configuration is the **Virtual-IO Manager** tab.

1. Launch the relevant wizard using the **Configuration Backup / Restore** button.
2. Select **Restore Configuration** to restore a configuration from a file.

3. Click **Next** to go to the **Select File** step.
4. Select the backup file from which you want to restore the configuration:
   - Select **Restore file from Management Station** to select a backup file from the management station.
   - Select **Restore local file** to select a file on the computer on which the GUI runs
5. When you click **Restore file from Management Station**, you select a backup file from the list on the management station.
6. When you click **Restore local file**, you select a backup file on the local computer on which the GUI runs. By clicking **Browse** the file selection dialog box opens, which allows you to select the relevant backup file. You can also enter the backup file name directly in the input field.
7. Once you have selected a backup file, click **Next** to go to the **Select Data** step to define the data you wish to restore.
8. Select the configuration you wish to restore under **Select configuration to restore**:
   - Select **Profiles** to restore the server profiles.
   - Select a server to restore the configuration saved for a blade server.

11.2.1 Restoring server profiles

1. Select **Profiles** in the **Select Data** step to restore server profiles. The **Restore Options** dialog box opens in which you specify additional parameters. You can also open this dialog box by clicking the **Options...** button in the **Profiles** row.
2. Select the server profiles which will be restored:

   If you select the **Restore only reassigned profiles**, only the profiles that are reassigned to the selected servers will be restored (see section "Restoring blade server configurations" on page 306). If you do not select the **Restore only reassigned profiles** option, all the profiles saved in the backup file will be restored.

3. Select the **Keep existing profiles** option, to keep the existing profiles. This option is selected automatically if you select the **Restore only reassigned profiles** option. If you do not select this option, all existing profiles are deleted before the configuration is restored.
4. If you select the **Keep existing profiles** option, select what is to happen with the existing profiles of the same name in **On Existing Profiles:**
   - The restore operation is canceled and an error message displayed (**Abort Restore**).
   - The existing profiles are renamed by adding the backup date and possibly also a number to the file names (**Rename Restored Profiles**).
   - The existing profiles are replaced by the profiles contained in the backup (**Replace Existing Profiles**).
   - The profiles are not restored and the existing profiles continue to be used instead (**Use Existing Profiles**).
5. If you want to assign new virtual addresses, select **Assign new virtual addresses**. The restored profiles are assigned new virtual addresses.
6. If you want to restore address ranges, select **Restore Address Ranges**.

   To restore address ranges will be necessary if address ranges have been changed since the backup configuration was saved or if they have been created on another management station.

   This option cannot be selected together with **Keep existing profiles**.
7. Click **OK** to apply your selection or **Cancel** to close the dialog box without applying your selection.

11.2.2 Restoring blade server configurations

1. To restore blade server configurations, select the corresponding blade server in the **Select Data** step. The **Restore Options** dialog box opens in which you specify additional parameters. You can also open this step

by clicking the **Options...** button in the table row.

2. Select **Restore Spare Information** to restore the information on the spare slots.
3. Select **Reassign Profiles** if you want to reassign the profiles which were assigned when the backup was performed. This option is only available if the backup contains profiles. If the option to restore profiles is not selected, this is performed automatically when you select the **Restore only reassigned profiles** option.
4. Select **Restore On Different Chassis** to restore the backup on another blade server. If you select this option, you have to select the destination blade server in another dialog box.
5. Click **OK** to apply your selection and **Cancel** to close the dialog box without applying your selection.

11.2.3 Restoring PRIMERGY rack server configurations

1. To restore PRIMERGY rack server configurations, select the corresponding PRIMERGY rack server in the **Select Data** step.

   The **Reassign Profiles** option is automatically selected and cannot be deselected. This means that the profiles which were assigned when the backup was performed are reassigned. You cannot specify more options for restoring a PRIMERGY rack server. If **Profiles** is not selected in the **Select Data** step of the **Restore Configuration** wizard, it is selected automatically when you select a PRIMERGY rack server. In this case, the **Restore only reassigned profiles** option will be selected for profiles (see section "Restore Options dialog box (server profiles)" on page 243).

11.3 Deleting backup files on the management station

The starting point for deleting backup files on the management station is the **Virtual-IO Manager** tab.

1. Launch the relevant wizard using the **Configuration Backup / Restore** button.
2. Select **Delete Configuration** to delete backup files you no longer need on the central management station. The second step **Select File** in the **Delete Backup Files** wizards opens.
3. Select the files you want to delete on the management station. To do this, select the corresponding files in the list. Using the **Select All** and **Clear Selection** buttons, you can select or deselect all the files.
4. Click **Next** to review the files to be deleted, and delete them by clicking the **Finish** button.

11.4 Restoring VIOM-specific configurations

11.4.1 Restoring an IBP module configuration

If an error occurs during the configuration of an IBP module, then the configuration of the IBP module no longer matches the one saved internally in the database in the Virtual-IO Manager.

This incorrect status is displayed on the **Setup** tab for the respective chassis with a corresponding status icon on the graphical display of the IBP module on the rear side of the chassis.

In this situation, you can only continue configuring the respective IBP module once you have restored the configuration. You do this as follows:

1. Click the respective IBP module in the graphical display on the rear of the chassis.
2. Then click the **Restore IBP** button.

The **Restore IBP** function rewrites the configuration saved internally in the VIOM database to the IBP module.

In doing so, all the connections are interrupted momentarily, as all the network connections are reprogrammed in the IBP module.

The same applies if a defective IBP module is replaced. In this case, the Virtual-IO Manager recognizes that a new connection blade has been inserted in the managed blade server chassis. The corresponding module is indicated on the **Setup** tab with a note regarding the inconsistent configuration.

You restore the configuration as follows:

1. Carry out the base configuration (configure the IP parameters, system name of the IBP, …) of the connection blade. In addition, the IP configuration of the connection blade must have been read from the ServerView Operations Manager. This is done as part of a regular scan of the blade server chassis.

   You may have to start the blade server chassis scan explicitly by selecting **Explore** in the context menu of the relevant blade server chassis in the server list of the ServerView Operations Manager.

2. Click **Restore IBP** on the **Setup** tab to write the configuration saved internally to the new connection blade.

11.4.2 Deleting the configuration of an uninstalled IBP module

If an IBP module for which a configuration still exists has been removed from a chassis, the configuration remains in the internal database of the Virtual-IO Manager.

This incorrect state is displayed on the **Setup** tab for the respective chassis with a corresponding state icon on the graphical display of the empty slot on the rear of the chassis. **slot is empty** is indicated as the state in **State Cause** for this slot.

You can delete the configuration you no longer need. You do this as follows:

1. Click the relevant slot in the graphical display on the rear of the chassis.
2. Then click the **Delete Configuration** button.

Delete Configuration deletes the configuration saved internally in the VIOM database.

11.4.3 Restoring the configuration of a server blade slot

In a few rare cases (e. g. when the network connection to the management blade is interrupted), an error may occur during the assignment of a server profile to a server blade slot. In this case, the following errors may occur:

- The configuration of one or several IBP modules is not correct.

  In this case, proceed as described in section "Restoring an IBP module configuration" (see "Restoring VIOM-specific configurations" on page 308).

- The virtualization data (virtual I/O addresses and possibly also boot parameters) of the profile have not been written correctly. In this case, the corresponding slot on the **Setup** tab indicates the status **Fault**.

  Here, you can use the **Restore Slot** button on the **Setup** tab to rewrite the virtualization data.

  If a server blade is in the slot, the server blade must be disabled for a **Restore Slot**.

Replacing the front control blade

If the front control blade of a blade server needs to be replaced, the configuration has to be restored for each slot. This means you have to carry out **Restore Slot** for each server blade slot.

In this case, VIOM also provides an option that allows you to restore the entire chassis (see section "Restoring the blade server chassis configuration" on page 311).

- The Virtual-IO Manager does not automatically recognize that the front control board has been replaced and hence all the related virtualization data has been lost.
- Before you restore the configuration of the slots or the entire configuration of the blade server chassis after replacing the front control board, you first have to carry out the base configuration of the management blade.

11.4.4 Restoring the blade server chassis configuration

By clicking **Restore** on the **Setup** tab, the entire virtualization information of the Virtual-IO Manager is rewritten for a blade server chassis managed by the Virtual-IO Manager:

- The configurations saved in the Virtual-IO Manager are performed again accordingly for all the IBP modules in the chassis.
- The virtualization data of the assigned profiles is saved again in the chassis using the existing configuration data in the Virtual-IO Manager for each server blade slot.

Any installed server blades must be disabled for this.

12 Importing and exporting server profiles

You have the option of exporting and importing server profiles. The exported data is available as XML files that can be modified or enhanced. You can use this procedure, for example, if you want to create a number of profiles without using the wizard. The format of the exported files or the files to be imported is described in section "Format of export files" on page 314.

12.1 Exporting server profiles

To export server profiles, proceed as follows:

1. In the left section of the **ServerView Virtual-IO Manager** window, switch to the **Server Profiles** view. If applicable, click **Profiles**.
2. Select the required profiles.
3. Then click the **Export** button.

   A file selection box opens in which you select the name of the file to which you want to save the exported profiles.

12.2 Importing server profiles

To import server profiles, proceed as follows:

1. In the left section of the **ServerView Virtual-IO Manager** window, switch to the **Server Profiles** view. If applicable, click **Profiles**.
2. Then click the **Import** button.

   A file selection box opens.
3. In the file selection box, select the file that you want to import.

12.3 Format of export files

The file specified during an import must be an XML file whose root element is **Objects** (see below). This file has the following structure:

```
<?xml version="1.0" encoding="UTF-8"?>

<Objects xmlns="http://schemas.fujitsu.com/serverview/viom/objects"

xmlns:xsi=http://www.w3.org/2001/XMLSchema-instance

schemaVersion="V3.1">
…

</Objects>
```

12.3.1 The Objects element

The **Objects** element is the root element of the XML file and must have the following attributes:

- xmlns="http://schemas.fujitsu.com/serverview/viom/objects"
- xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
- schemaVersion="V3.1"

This element can also have the following optional attribute:

errorAction

Specifies whether the import will continue if an error occurs. The value can be overwritten in lower-level elements. Possible values are:

Abort

Default

The first error aborts the entire import.

Continue

Errors are converted into warnings. The import then continues with the next object.

The **Objects** element contains the following elements:

- **UserInfo** (optional)
- **ServerProfiles** Describes the profiles to be imported (see below)

12.3.2 The ServerProfiles element

The **ServerProfiles** element can have the following optional attributes:

errorAction
Specifies whether the import will continue if an error occurs. If this attribute is specified, it overwrites the value from the **Objects** element. For individual server profiles, the value in the corresponding **ServerProfile** element can be overwritten. Possible values are:

Abort
The first error aborts the entire import.

Continue
Errors are converted into warnings and the import then continues with the next object.

existingProfileAction
Specifies the response if a server profile to be imported already exists. For individual server profiles, the value in the corresponding **ServerProfile** element can be overwritten. Possible values are:

Refuse
Default

Handled in accordance with the current error behavior (see **errorAction**).

Replace
An existing server profile is replaced if it is not assigned to a slot. Otherwise, the procedure is the same as above (see **Refuse**).

existingAddressAction

Specifies the response if a virtual address is already in use. For individual server profiles, the value in the corresponding **ServerProfile** element can be overwritten. Possible values are:

Refuse

Handled in accordance with the current error behavior (see **errorAction**).

New

A new address is allocated for each affected IO channel.

NewForAll

New addresses are assigned for all IO channels of the profiles.

It contains a **ServerProfile** element for each server profile to be imported (see below).

12.3.3 The ServerProfile element

Each **ServerProfile** element contains the following attributes:

id

Is the server profile ID. This value must be unique within the import file. It is not written to the VIOM database, but it is specified in error messages.

errorAction

optional

Specifies whether the import will continue if an error occurs. If this attribute is specified, it overwrites the value from the **ServerProfiles** or **Objects** element. Possible values are:

Abort

The first error aborts the entire import.

Continue
Errors are converted into warnings and the import then continues with the next object.

existingProfileAction
optional

Specifies the response if a server profile to be imported already exists. If this attribute is specified, it overwrites the value from the **ServerProfiles** element. Possible values are:

Refuse
Handled in accordance with the current error behavior (see **errorAction**).

Replace
The existing server profile is replaced if it is not assigned. Otherwise, the procedure is the same as above (see **Refuse**).

existingAddressAction
optional

Specifies the response if a virtual address is already in use. If this attribute is specified, it overwrites the value from the **ServerProfiles** element. Possible values are:

Refuse
Handled in accordance with the current error behavior (see **errorAction**).

New
A new address is allocated for each affected IO channel.

NewForAll
New addresses are assigned for all IO channels of the profile.

It contains the following elements:

ServerProfileName
Name of the server profile.

IOVirtualizationUsage

Yes
The profile uses virtual addresses. You must specify the **Address-Virtualization** element in all **IOChannel** elements of this profile.

No
The profile does not use any virtual addresses. You must not specify the **AddressVirtualization** element in any **IOChannel** element of this profile.

BootMenuUsage

Yes
The VIOM boot settings can be overwritten on your local computer.

No
The VIOM boot settings are prevented from being overwritten on your local computer.

Comment
optional

Comment on a more detailed description of the profile.

SmuxSettingMezzanine2
optional

If a second mezzanine card is defined as a LAN card, this element is used to specify the fabric to which the card is routed. Possible values are:

Fabric3
All paths are routed to fabric 3.

Fabric4
Default

All paths are routed to fabric 4.

Fabric3+4
LAN1 is routed to fabric 3 and LAN2 to fabric 4.

IOChannels
Defines the ports. This element contains an **IOChannel** element for each port of the profile (see below).

12.3.4 The IOChannel element

An **IOChannel** element contains the following elements:

IOChannelSpec
Specifies the port and contains the following elements:

IOChannelType
Specifies the port type:

LAN
LAN port

FC
FC port

ISCSI
For future use

LANFunction
Physical LAN function of a CNA port

FCFunction
Physical FCoE function of a CNA port

ISCSIFunction
Physical iSCSI function of a CNA port

IOSlotIndex
For I/O Channels in a profile that uses only one slot:always 0. For multi-slot profiles: 0 for the main slot and 2 for the lower slot.

IOBoardType
Type of board to which the port belongs:

OnBoard
OnBoard

DaughterCard
Mezzanine card

AddonCard
PCI card

IOBoardNumber
Number of the board to which the port belongs: always 1 for OnBoard, 1 or 2 for mezzanine cards, and 1 to 14 for PCI cards.

IOPortNumber
Number of the port, counts upwards from 1.

IOFunctionNumber
Number of the physical function of a CNA port. If this **IOChannel** element do not specify a physical function, the value must be 1.

IOChannelUsage
Indicates whether the port can be used.

Yes
The IO channel is "enabled".

No
The IO channel is "disabled".

Networks

Defines the networks used by this IO channel. This element contains an optional **NetworkName** element optionally followed by one or more **ServiceName** elements and then optionally by one or more **TaggedName** elements:

NetworkName

Optional, must not be specified for FC ports or for PRIMERGY rack server profiles.

Network name. If you want to use the profile on blade servers with IBP modules, you can specify a network. If you work with blade servers that have non-VIOM-capable LAN modules (Open Fabric mode), do not specify a network as it is not possible to define networks on these modules.

ServiceName

Optional, must not be specified for FC ports or for PRIMERGY rack server profiles.

Service network name.

TaggedName

Optional, must not be specified for FC ports or for PRIMERGY rack server profiles.

Tagged VLAN network name.

The **NetworkName, ServiceName**, and **TaggedName** elements must be the same for all physical functions of a CNA port.

AddressVirtualization

(Must not exist if the profile does not use any virtual addresses; must exist if the profile uses virtual addresses).

Contains an **Address** element for each virtual address (see "The Address element" on page 323).

BootDeviceUsage
Indicates whether the port is used as a boot device.

Yes
The port is a boot device and the **BootEnvironment** element (see below) must be specified.

No
The port is not a boot device and the **BootEnvironment** element (see below) must not be specified.

BootEnvironment
optional

Specifies the boot settings for this port; it is described below. This element must be specified if **BootDeviceUsage** has the value **Yes** (see above).

DCBUsage
Indicates whether DCB (Data Center Bridging) should be used for this port or function. It should only be enabled for FCoE functions.

Yes
DCB will be used and the **DCBConfiguration** element (see below) must be specified.

No
DCB will not be used and the **DCBConfiguration** element (see below) must not be specified.

DCBConfiguration
optional

Specifies the DCB configuration for this function. This element must be specified if **DCBUsage** has the value **Yes** (see above).

FunctionConfiguration
optional

Specifies the configuration for a function (see " The FunctionConfiguration element" on page 329). This element must be specified for functions and must not be specified for other **IOChannelType**.

12.3.5 The Address element

An **Address** element that specifies the E-MAC address of a CNA FCoE-function must have the attribute **purpose="fcoe"**. Each other **Address** element must have the attribute **purpose="normal"**.

If it does not contain any elements, the VIOM Manager automatically assigns the virtual addresses according to the port type during the import.

A LAN port, CNA LAN-function, or CNA iSCSI-function may contain the **VirtualMAC** element whose value is the virtual MAC address.

An FC port may contain the **VirtualWWNN** and/or **VirtualWWPN** elements whose values are the corresponding virtual addresses. If only one of these two elements is specified, the VIOM Manager automatically assigns the other virtual address during the import.

A CNA FCoE-function may contain the **VirtualMAC** element additional to the elements of a FC port. If it is not specified, the VIOM Manager assigns the E-MAC address during import.

The **ignoreRange** attribute of the **VirtualMAC**, **VirtualWWNN**, and **VirtualWWPN** elements specifies whether the given virtual address is in the range that was specified when you installed VIOM. Possible values for this attribute are **Yes** (ignore the range) and **No** (check the range).

12.3.6 The BootEnvironment element

The **BootEnvironment** element contains the **BootPriority** element and one of the following elements:

- **PXEBootConfiguration**
- **ISCSIBootConfiguration**

- **FCBootConfiguration**

BootPriority
Specifies the boot order. Possible values are 1 to 4. However, each value can only appear once within a profile.

PXEBootConfiguration
Specifies that the port is a PXE boot device. This element is empty. You can only specify this element for LAN ports.

ISCSIBootConfiguration
Specifies that the port is an iSCSI boot device. You can only specify this element for LAN ports. The elements contained in this element are described below.

FCBootConfiguration
Specifies that the port is a SAN boot device. You can only specify this element for FC ports. The elements contained in this element are described below.

12.3.7 The ISCSIBootConfiguration element

The **ISCSIBootConfiguration** element contains the following elements:

ISCSIInitiator
Specifies the values of the iSCSI initiator. This element contains the following elements:

DHCPUsage

Yes
In the case of an iSCSI boot, the system tries to obtain the client IP address, subnet mask, and gateway IP address from a DHCP server. Only the initiator name must be specified here.

No

A static client IP address, subnet mask, and gateway IP address must be specified.

Name

Name of the iSCSI initiator to be used (in the case of an iSCSI boot) for the connection to the iSCSI target.

VLANId

optional

VLAN ID that is used by the HBA to send its requests. It should be specified only for CNA-iSCSI-functions.

IPv4Address

Only when **DHCPUsage** is **No**

The static client IP address to be used for this port. The port will use this IP address for the entire iSCSI session.

SubnetMask

Only when **DHCPUsage** is **No**

The IP subnet mask. This should be the IP subnet mask of the network used to connect this port (in the case of an iSCSI boot).

GatewayIPv4Address

Only when **DHCPUsage** is **No**

The IP address of the network gateway. This is necessary if the iSCSI target is in a subnetwork other than the subnetwork of the selected iSCSI boot port. If no gateway is used, **0.0.0.0** can be specified as IP address of gateway.

ISCSITarget

Specifies the values of the iSCSI target. This element contains the following elements:

DHCPUsage

Yes
In the case of an iSCSI boot, the system tries to obtain the name of the iSCSI target, the IP address of the iSCSI target, the IP port number, and the SCSI LUN ID from a DHCP server in the network.

No
You must specify a static name for the iSCSI target, a static IP address for the iSCSI target, a static IP port number, and a static SCSI LUN ID.

Name
Only when **DHCPUsage** is **No**

The IQN name of the iSCSI target.

IPv4Address
Only when **DHCPUsage** is **No**

The IP address of the iSCSI target.

PortNumber
Optional; only if **DHCPUsage** is **No**

TCP port number (default: 3260 for iSCSI).

BootLUN
Only when **DHCPUsage** is **No**

The LUN ID of the boot disk on the SCSI target

AuthenticationMethod

None
No authentication is used.

CHAP
CHAP authentication is activated for this port. CHAP allows the target to authenticate the initiator.

MutualCHAP
Mutual CHAP authentication is activated for the port. Mutual CHAP allows the initiator to authenticate the target.

ChapUserName
Not when **AuthenticationMethod** is **None**

The CHAP user name. The name must be identical to the name configured on the iSCSI target.

ChapSecret
Not when **AuthenticationMethod** is **None**

The CHAP password. This password must be identical to the password configured on the iSCSI target and it must contain 12 to 16 characters. This password must differ from the password in the **MutualChapSecret** element.

MutualChapSecret
Only when **AuthenticationMethod** is **MutualCHAP**

The Mutual CHAP password. This password must be identical to the password configured on the iSCSI target and it must contain 12 to 16 characters. This password must differ from the password in the **ChapSecret** element.

12.3.8 The FCBootConfiguration element

The **FCBootConfiguration** element contains the following elements:

FCTarget
Once or twice

The first or only **FCTarget** configures the boot device while the second FCTarget configures the backup boot device. It contains the following elements:

TargetWWPN
The WWPN (worldwide port name) of the port for the boot device.

TargetLUN
The LUN (logical unit number) address of the boot device.

FCLinkSpeed
Specifies the transmission speed used by this port. You can specify the following values:

0
auto negotiate

The transmission speed is negotiated with the external switch.

1
1 Gbit/s full-duplex

2
2 Gbit/s full-duplex

4
4 Gbit/s full-duplex

8
8 Gbit/s full-duplex

FCTopology
Specifies the type of port connection with the external SAN network. You can specify the following values:

0
auto (loop first)

4
Point-to-Point

8
auto (Point-to-Point first)

12
Arbitrated loop

12.3.9 The DCBConfiguration element

The **DCBConfiguration** element specifies one priority group and contains the following elements:

PriorityPurpose
The purpose of this priority group.

Possible values are **Other**, **FCoE**, and **iSCSI**.

PriorityLevel
Priority level; default for FCoE is 3, default for iSCSI is 4.

Possible values are 0 to 7.

12.3.10 The FunctionConfiguration element

The **FunctionConfiguration** element specifies function-specific attributes:

FunctionBandwidth
The share of the bandwidth in percent that is reserved for this function. If the sum of all bandwidths of one IO-channel is not 100, the values are internally adjusted accordingly.

FunctionVLANId
optional

The Vlan ID that is used by this function. This value must not be specified for FCoE functions.

13 VIOM scenarios

This chapter describes examples of when VIOM can be used.

13.1 Shifting tasks from one server blade to another

If an operating system or an application needs to run on another server blade of a blade server or if a server blade fails and another server blade has to assume these tasks, then VIOM allows server profiles to be moved from one slot of the blade server to another from the central management station.

By virtualizing the I/O addressing, tasks can be shifted without the network administrators having to be involved.

In order to shift server profile tasks from the central management station without involving a network administrator, you must use virtual addresses. To do this, you have to select virtual addressing when you define the server profile.

1. When defining a server profile, check the **Use virtual addresses** option in the second step of the **Create Server Profile** wizard.

You move a server profile from one server blade to another on the **Server Configuration** tab. You do this as follows:

1. Click **Server List**, if applicable, to switch to the server list view in the area on the left.
2. In the tree structure on the left, select the corresponding blade server from the **VIOM Managed** group.
3. Switch to the **Server Configuration** tab in the area on the right.
4. Select the server blade in the table that is to assume the tasks of the faulty server blade.

5. You can only assign a server profile to a server blade if the server blade is shut down.
   Therefore shut the server blade down, if applicable, by clicking **Shutdown** or by selecting **Shutdown** in the context menu. You can see the state of a server blade in the **State** column. To see the current power state, you may need to update the display using the **Update States** button or with **Update State** in the context menu.
6. Click **Assign Profile** or select **Assign Profile** in the context menu of the required server blade.The **Select Profile** dialog box opens (see "Select Profile dialog box" on page 245.
7. Select the server profile of the faulty server blade in the tree structure. The area on the right displays information on the selected profile. Confirm your selection with **OK**.
   A message appears asking if you wish to delete the assignment. If you confirm this, then the previous assignment is deleted and the profile is assigned to the new server blade.
   The display on the **Server Configuration** tab is updated accordingly.
8. Then switch the server blade back on using the **Boot** button.

The new server blade can then assume the tasks of the previous one without any restrictions.

13.2 Moving tasks using the server profile failover

Tasks can also be moved from one server blade to another. To do this, you define spare slots. It is advisable to install server blades in the spare slots that you do not use for regular operations.

You can also use a spare slot as a completely normal slot even though it has been configured as a spare slot.

1. Click **Server List**, if applicable, to switch to the server list view in the area on the left.
2. In the tree structure on the left, select the corresponding blade server from the **VIOM Managed** group.
3. Switch to the **Server Configuration** tab in the area on the right.

4. Define the slots that are to be available as spare slots for a server profile failover. To do this, enable the check box in the **Spare** column in the table for the relevant slots.

If an error occurs or maintenance work needs to be carried out, select the Failover function in the context menu of the affected server blade on the **Server Configuration** tab. In this case, VIOM searches for the slot configured as a spare slot. When it has found a suitable slot, it deletes the previous profile assignment and assigns the profile to the spare slot. The server blade installed in the spare slot therefore assumes the role of the failed server blade including the network addresses.

- The profile can only be assigned to the spare slot if the server blade in this slot is switched off (power off) (see also section "Assigning server profiles" on page 296).
- The "spare" blade only assumes the role of the failed server blade if the VIOM server profile also contains the boot settings (e.g. FC boot settings) and if the spare blade is suitable for booting the same operating system as booted by the original blade. The Virtual-IO Manager does not make any changes to the operating system.

13.3 Disaster Recovery

In a disaster recovery scenario - meaning you have two generally identical sites and if one fails you restart all or at least the most important applications on the other site - the server profile assignment feature may be used to “move the application” from one site to the other as illustrated in the following diagrams:

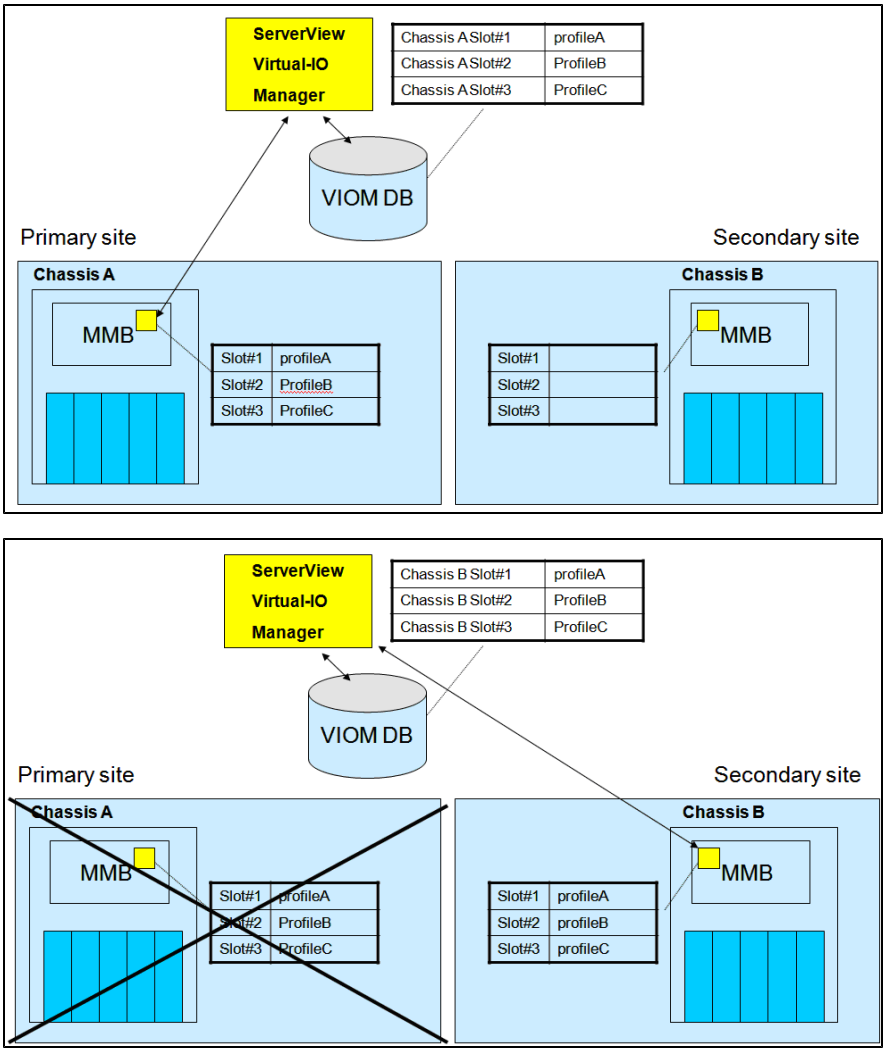

This means you shut down the blade on the failed site, unassign the profile, and assign it again on the other site. But this only works if the chassis (MMB, connection blades, and server blades) are accessible to the VIOM management station, which is typically not the case in such a disaster scenario.

To allow the switchover in this case, VIOM provides a “forced unassign” option, allowing you to unassign a profile even if the corresponding MMB (containing this profile) is not available.

But applying this function means that there will be two chassis with at least partially identical profiles, potentially resulting in, for example, duplicate virtual addresses as soon as the failed site goes online again.

This chapter describes the settings and procedures for avoiding duplicate addresses and how to handle a site failover with VIOM.

As already mentioned, there are two major issues: unassign the profiles at a failed / no longer reachable site, and handle the power-on procedure of the failed site to avoid duplicate MAC and WWN addresses.

The first issue is solved by the “forced unassign” function, which is available as a menu item in the VIOM Web GUI as a CLI command.

But beware: This function should only be used if there is no other way to get the profile unassigned! A forced unassign will only remove the “profile-to-server blade assignment” in the VIOM database but not in the corresponding hardware (MMB, IBP, server blade). If you later assign this profile to a different server blade you may end up with two different blades containing identical virtual addresses.

Therefore you should only assign and boot the profile on the other site if you are sure that the failed site is not longer active on the networks. The easiest way to ensure this is to completely power-off the failed site before restarting the profiles on the other site.

To avoid having two blades powered-on with the same server profile, you must prepare the power-on behavior of the chassis. During site preparation the server blades’ **BehaviorAfterACFail** flag need to be set to **Always power off**. You can do this in either of the following ways:

Via MMB: (preferred because more convenient)

The MMB “Power Restore Policy” function on the **Power Management** tab controls the “Behavior After AC Fail”. Set it as indicated to **Always power off** for every server blade in the chassis.

Via iRMC on each server blade:

As an alternative you may log on to each iRMC and set the “Power Restore Policy” to “Always power off” for each iRMC separately, at least for all server blades potentially running virtual addresses. We strongly recommend that you set it for all.

This setting prevents a failed site from automatically booting the server blade after AC power becomes available again.

Procedure for restoring the failed site

After powering-on the failed chassis, the blades will automatically execute a so-called inventory boot and will be switched off again automatically.

As described above, the VIOM database contains the current profile assignments, but the failed MMB still contains the status prior to the disaster. Synchronizing the hardware with the VIOM database may be achieved by either restoring the whole chassis, which is preferable because it is simple and secure.

By choosing **Restore** on the **Setup** tab, you are ensuring that all chassis components - MMB, connection blades, server blades, and IBP - are correctly set, but this only works if all the server blades are switched off.

Alternatively, individual server blades (slots) may be restored by applying the **Restore Slot** function, but keep in mind that this only configures the selected slot and does not restore the entire IBP. As an advantage this function works in case if certain server blades need to be kept running.

After applying the restore function, the VIOM database and the chassis hardware of the restored site are in sync again and you may move applications back to the recovered site using the server profile assignment function.

14 VIOM database

This chapter describes the VIOM database on a Windows- or Linux-based CMS.

The Virtual-IO Manager requires two databases:

- The ServerView database, which is installed with ServerView Operations Manager, and
- The VIOM database, which is installed with VIOM.

Both databases should be regularly backed up. The ServerView database is backed up independently of the VIOM database.

- For backing up the VIOM database you can use the VIOM Backup Service, which is described in this chapter (see section "VIOM Backup Service" on page 338).
- For the ServerView database there is a separate backup concept, which is described in "Installing ServerView Operations Manager Software under Windows - Installation Guide" and "Installing ServerView Operations Manager Software under Linux - Installation Guide".

On Windows: In addition to the database backups, transaction logs are created for both the ServerView database and the VIOM database. In the event of a problem, the database can be restored using the most recent backups of the ServerView and VIOM databases and their respective transaction logs.

Note that under Linux and PostgreSQL no transaction logs are saved - only full backups.

- For how to restore the VIOM database on Windows using the backups and transaction logs, see section "Restoring the VIOM database on Windows" on page 345.
- For how to restore the VIOM database on Linux using the backups, see section "Restoring the VIOM database on Windows" on page 345.
- For how to restore the ServerView database , see "Installing ServerView Operations Manager Software under Windows - Installation Guide" and

"Installing ServerView Operations Manager Software under Linux - Installation Guide".

14.1 VIOM Backup Service

The Backup Service of the Virtual-IO Manager is a service used to periodically back up the VIOM database.

Note that its sole purpose is currently to back up the VIOM database, whereas the backup of ServerView database is done by ServerView itself. But a backup of the VIOM database will be of little use if the ServerView database is lost. So, in general, you should back up both databases periodically at roughly the same times.

Under Windows, the Backup Service is designed as a Windows service using the Open Source Quartz framework.

The service schedules backup jobs for the databases currently used by VIOM, namely SQL Server on Windows respectively PostgreSQL on Linux.

The Backup Service allows three types of backup, for which there is one backup job each:

- A full backup of the database (**full backup job**), on Windows and Linux
- An incremental backup of the database (**incremental backup job**), on Windows only
- The backup of the transaction logs (**transaction backup job**), on Windows only

Note that under Linux and PostgreSQL no transaction logs are saved - only full backups.

You may modify the schedule of these backup jobs in a syntax similar to UNIX cron format.

The Backup Service is integrated in the install packages of VIOM Manager. The Backup Service does not start directly after installation, as it must be configured first. After this you must start it manually.

The following parameters need to be configured for the backup jobs:

- The Quartz cron expressions which determine when the backup jobs are scheduled (see section "Configuring the job schedule on Windows" on page 339) and "Configuring the job schedule on Linux" on page 342.
- The parameter for the directory where the output files are stored (see section "Configuring the output directories" on page 343)

Dependencies on the database

On Windows: Both the VIOM Manager and the Backup Service use Windows Authentication to access SQL Server. This means, at the very least, that both services must run under the same Windows account as the SQL Server.

On Linux: The Virtual IO Manager is installed under the root user and with root privileges. The database is accessed via the database user **svuser** which has already been created by ServerView prior to the Virtual-IO Manager installation.

Warning

On Windows: It is a known behaviour of SQL Server that the transaction log grows unless it is periodically backed up. We therefore strongly recommend that you configure and start the Backup Service in order to backup at least the transaction log.

14.1.1 Configuring the job schedule on Windows

The jobs and their triggers are defined in the **quartz_jobs.xml** file in the **<installation_path>\VIOM Manager** directory. Cron triggers may be modified by the user by editing this XML file. The database configuration is taken from the **ViomConfig.properties** file.

Trigger description relating to the full backup job in the quartz_jobs.xml file

```
<trigger>
 <cron>
  <name>cronFullBackupJobTrigger</name>
  <group>MSSQL-cron</group>
  <description />
  <job-name>FullBackupJob</job-name>
  <jop-group>MSSQL</jop-group>
  <cron-expression>0 0 19 ? * FRI</cron-expression>
 </cron>
</trigger>
```

The trigger description above only belongs to the job called **FullBackupJob** of SQL Server. There are also trigger descriptions for the **IncrementalBackupJob** and **BackupLogJob** jobs.

To change the execution time for a backup job, you must open the **quartz_jobs.xml** file and update the Quartz cron expression in the trigger description relating to the relevant backup job. For a description of Quartz cron expressions, see section "Syntax of Quartz cron expressions" on page 340.

14.1.1.1 Syntax of Quartz cron expressions

This section provides a brief explanation of the Quartz cron expression syntax. In the context of the VIOM Backup Service, a cron expression describes when one of the database backup jobs defined for your database will execute.

A Quartz cron expression is a sequence of six to seven fields separated by a blank. The following fields are allowed:

Field	Mandatory	Values	Special characters
Seconds	yes	0 - 59	, - * /
Minutes	yes	0 - 59	, - * /
Hours	yes	0 - 23	, - * /
Day of Month	yes	1 - 31	, - * / ? L W

Field	Mandatory	Values	Special characters
Month	yes	1 - 12 or JAN - DEC	, - * /
Day of Week	yes	1 - 7 or SUN - SAT	, - * / ? L #
Year	no	1970 - 2099	, - * /

Table 5: Fields of a Quartz cron expression

Example

A simple example would be:

```
0 0 19 24 DEC ? 2010
```

This cron expression would trigger a backup job at 7.00 p.m. on 24.12.2010.

The following table shows the use of special characters in a cron expression:

Special characters	Description	Example	Meaning of the example
x,y	Describes a list of 2 or more values	* * 9 ? * SAT,SUN	At 9 a.m. every Saturday and Sunday
x-y	A range of values from x to y inclusive.	* * 6 ? * MON-FRI	At 6 a.m. every Monday to Friday
*	Describes all possible values	* * * ? * *	Every second
x/y	Describes a sequence starting with x and incremented by y	* 0/5 * 1 * ?	Every 5 seconds starting at 0:00 on the first day of each month

Special characters	Description	Example	Meaning of the example
?	Day of month and day of week must not be given a value at the same time, so use the question mark for the unspecified field of the two	* * * ? * MON 2010	Every second on all Mondays in 2010
L	Last day of month	0 0 8 L * ?	At 8.00 a.m. on the last day of every month
xL	Last weekday x of every month, where x ranges from 1-7 (meaning SUN-SAT)	0 0 12 ? * 6L	At 12 noon on the last Friday of every month
xW	Nearest weekday to x, where x ranges from 1 to 31	0 0 9 15W * ?	At 9 a.m. on the nearest weekday to the 15th of every month
x#y	The y-th weekday x of a month, where y ranges from 1 to 5 and x from 1 to 7 (meaning SUN-SAT)	0 0 9 ? * 2#1	At 9 a.m. on the first Monday of every month

Table 6: Use of special characters

14.1.2 Configuring the job schedule on Linux

The Backup Service uses the open source Quartz framework to periodically run jobs that perform various backups of the database. Quartz is configured by editing the file:

/opt/fujitsu/ServerViewSuite/plugins/viom/Manager/quartz_jobs.xml

For PostgreSQL:

If the schedule does not suit your needs, you should edit the Quartz cron expression:

```
<cron>
 <name>PostgresJobTrigger</name>
 <group>Postgres-cron</group>
 <description>…</description>
 <job-name>PostgresJob</job-name>
 <job-group>Postgres</job-group>
 <cron-expression>0 0 19 ? * FRI</cron-expres-
sion>
</cron>
```

14.1.3 Configuring the output directories

To configure the output directories for the backup jobs, you must open **quartz_job.xml** (e.g. on Windows in the **<installation_path>\VIOM Manager** directory).

Example

```
<?xml version='1.0' encoding='utf-8'?>
<!DOCTYPE quartz [
 <!ENTITY outputDirLinux "/var/f-
ujitsu/ServerViewSuite/viom/postgres/backups">
 <!ENTITY outputDirWindows "c:\Backups">
 <!ENTITY outputLogDirWindows "c:\Backups\Log">
]>
```

You must modify the XML entity definitions at the beginning of the file, so that they suit to your requirements:

- The **outputDirWindows** entity defines the directory for full and incremental backup files (example: **c:\Backups**) on Windows.
- The **outputLogDirWindows** entity defines the directory for the backup files of the transaction logs (example: **c:\Backups\Log**) on Windows.

- The **outputDirLinux** entity defines the directory for full backup files (example: **/var/fujitsu/ServerViewSuite/viom/postgres/backups**) on Linux.

The output directories for the backup files should be on a different hard disk than the VIOM database. This could also be an external hard disk.

14.1.4 Starting the Backup Service on Windows

When you have configured the Backup Service you must start the **ServerView Virtual IO DB Backup Service**.

1. Select **Start – [Settings] – Control Panel – Administrative Tools – Services**.
2. Select the **ServerView Virtual IO DB Backup Service** and then select **Start** from the context menu.

To have the Backup Service start automatically from now on, you will need to configure it as follows:

1. Select **Start – [Settings] – Control Panel – Administrative Tools – Services**.
2. Select the **ServerView Virtual IO DB Backup Service** and then select **Properties** from the context menu.
3. On the **General** tab, set the start type to **Automatic**.
4. Click **OK**.

14.1.5 Starting the VIOM Backup Service on Linux

When you have configured the Backup Service you must start the Backup Service.

1. `service viom_backup start`
2. `chkconfig viom_backup on`

The second command is necessary for the Backup Service to restart after a reboot of the management station.

14.1.6 Logging the Backup Service

The Backup Service logs important events in the Windows Event Logging or syslog on Linux.

So you can search the Windows Event Logging respectively syslog for information on the Backup Services or for troubleshooting them.

14.2 Restoring the VIOM database on Windows

In the event of an error, you can restore the VIOM database from the backups. You must not delete the current VIOM database or the errored VIOM database. If you find any errors during restoration, you must restart the restoration from the beginning.

For the restoration, you must first read in the database backup and then, if available, one or more transaction log backups.

Backups of transaction logs are available if the corresponding backup jobs are configured in the Backup Service.

Before you restore the VIOM database, the ServerView database may also have to be restored. For more on restoring the ServerView database, please see the manual "Installing ServerView Operations Manager Software under Windows - Installation Guide".

When the VIOM Backup Service is started, the master database is backed up once. Follow the same steps as described below to restore the master database as well, but note that ServerView Operations Manager also backs up the master database. It is up to the user to choose the most recent backup for that database and restore it.

14.2.1 Restoration via SQL Server Management Studio

To restore the VIOM database and, if available, the transaction logs, proceed as follows:

Stop the services **ServerView Virtual IO Manager Services** and **ServerView Virtual IO DB Backup Services**:

1. Select **Start – [Settings] – Control Panel – Administrative Tools – Services**.
2. Select the appropriate service and then select **Stop** from the context menu.

Restrict the access to the VIOM database:

1. Start SQL Server Management Studio.
2. Connect with the SQL Server instance and select **Databases – ViomDB**.
3. Select **Properties** from the context menu.
4. Select the **Options** page and, under **Restrict Access**, select the entry **RESTRICTED_USER**. Click **OK** followed by **YES**.

Restore from the database backup:

1. Click **Databases** and then select **Restore Databases …** from the context menu.

   The **Restore Databases** window opens:

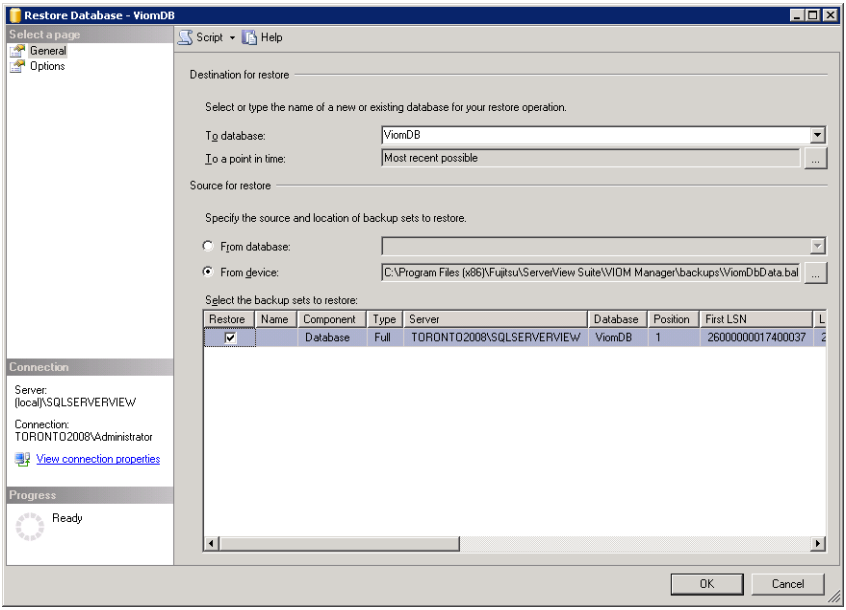

2. On the **General** page, enter the name **ViomDB** or **master** in the **To database** field or select the name from the list.
3. Select the option **From device**.
4. Click the **...** button.
5. Add the database backup **ViomDBData.bak** from the appropriate storage location and then click **OK**.
6. Select the database backup to be restored:
   In **Select the backup sets to restore**, click the box in the **Restore** column.
7. Switch to the **Options** page.
8. Select the options **Overwrite the existing database**.
9. If no transaction log backup **ViomDBLog.bak** is available, click **OK**. Otherwise, select the option **Overwrite the existing database and Leave database nonoperational and do not roll back uncommitted transactions. Additional transaction logs can be restored**.

10. Click **OK** and then **OK** again.

Restoration of the database begins. In the ObjectExplorer you will see the message **ViomDB (Restoring…)**.

Restore the transaction logs if available:

1. Click **Databases** and then select **Restore Databases …** from the context menu. The **Restore Database** window opens.
2. On the **General** page, in the **To Database** field select the name **ViomDB** from the list.
3. Select the option **From Device**.
4. Click the **...** button.
5. Add the transaction log file **ViomDBLog.bak** from the appropriate storage location and then click **OK**.
6. Select the database backup to be restored:

   In **Select the backup sets to restore**, click the box in the **Restore** column.
7. Switch to the **Options** page.
8. Select the option **Overwrite the existing database**.

   If you want to restore further transaction logs, select the option **Leave the database non-operational and do not roll back uncommitted transactions. Additional transaction logs can be restored.**

   If you want to restore the last transaction log, select the option **Leave the database ready to use ... Additional transaction logs cannot be restored**.
9. Click **OK** and then **OK** again.

Once the last transaction log has been restored, the database status must be normal again. The add-on **(Restoring…)** is no longer displayed in the **ObjectExplorer**.

The services **ServerView Virtual IO Manager Services** and **ServerView Virtual IO DB Backup Services** must be restarted:

1. Select **Start – [Settings] – Control Panel – Administrative Tools – Services**.
2. Select the appropriate service and then select **Restart** from the context menu.

It is now possible to access the VIOM database via Virtual-IO Manager again.

14.2.2 Restoration via Enterprise Manager

Enterprise Manager may be used for SQL Server 2000 to restore the VIOM database. How this is done for the ServerView database is described in the manual "ServerView Operations Manager - Installations under Windows". The restoration of the VIOM database is similar.

14.2.3 Checking the database backup

From time to time you can check the backups with the SQL Server Management Studio or the Enterprise Manager.

For how to check database backups with the Enterprise Manager, see the relevant sections in the manual "ServerView Operations Manager - Installations under Windows".

For the SQL Server Management Studio, follow the instructions in thesection "Restoration via SQL Server Management Studio" on page 345 but with the following changes:

1. On the **General** tab, enter any name in the **Restore as database** field, e. g. **RECOVERYTEST**.
2. On the **Options** tab, change the path names in the **Move to physical file name** column as follows:
   - **ViomDB.mdf** to **RecoveryViomDB.mdf**
   - **ViomDB_log.LDF** to **RecoveryViom_log.LDF**

Make all other entries as described in the relevant sections. Afterwards, the database should have been restored under the name **RECOVERYTEST**. You can check this as follows:

1. Click the SQL Server instance and select **Databases**.
2. Select **Refresh** from the context menu.

The database **RECOVERYTEST** must be displayed in the list.

You can then delete the database **RECOVERYTEST** as follows:

1. Select the database and then select **Drop** from the context menu.

14.3 Restoring the VIOM database on Linux

In the event of an error, you can restore the VIOM database from the backups. You must not delete the current VIOM database or the errored VIOM database. If you find any errors during restoration, you must restart the restoration from the beginning.

Before you restore the VIOM database, the ServerView database may also have to be restored. For more on restoring the ServerView database, please see "Installing ServerView Operations Manager Software under Linux - Installation Guide".

Here are the additional steps for restoring the VIOM database:

1. Stop the VIOM Backup Service:

   ```
   /etc/init.d/viom_backup stop
   ```

2. Stop the VIOM Manager Service:

   ```
   /etc/init.d/viom_man stop
   ```

3. Import the VIOM database:

   ```
   bzip2 -cd /var/fujitsu/ServerViewSuite/viom/postgres/
   backups/ViomDB.dump.bz2 |/opt/fujitsu/ServerViewSuite/
   Postgresql/pgsql/bin/psql -p 9212 -d ViomDB -U svuser
   ```

4. Start the VIOM Manager Service:

   ```
   /etc/init.d/viom_man start
   ```

5. Start the VIOM Backup Service:

   /**etc/init.d/viom_backup start**

15 Appendix

15.1 Replacing IBP modules

If an IBP connection blade in a chassis managed by Virtual-IO Manager fails, please perform the following actions when replacing this connection blade with a connection blade of the same type:

1. Unplug all LAN cables connected to the defective IBP.

   Important

   All cables should have a label that uniquely identifies them and allows you to reconnect them to the same ports in the new IBP module at the end of the procedure. The labeling should also contain information about active ports and backup ports. Before removing the LAN cables, please make sure you have all the information you need to reconnect them to the same ports on the new connection blade.

2. Now remove the defective connection blade and replace it with a new one of the same type.

   Important

   Do not connect the LAN cables at this stage!

3. After booting the new connection blade, check its mode. If it is not running in IBP mode, please change the mode to IBP mode and reboot the connection blade. You can change the mode via the web-based user interface of the management blade; see **Configuration** tab for the selected connection blade. Check and change the **Firmware Mode Setting** in the MMB web-based user interface.

4. Make sure that the IP configuration of the new connection blades is correct and also that the new connection blade has the correct authentication data (same user name and password) as the old one. If you are

using telnet as the communication protocol, please check whether telnet is enabled on the new connection blade.

5. When the new connection blade has the correct setup and is running in IBP mode (final boot has finished and connection blade is accepting IBP commands), please perform the action **Explore** in ServerView Operations Manager (menu item of the context menu in the **ServerList** window) for the related blade server chassis and wait until this action is complete.
6. Now start the user interface of Virtual-IO Manager and select the corresponding blade server chassis in the server tree (sub-tree **VIOM Managed**). On the **Setup** tab select the new connection blade. VIOM should report the error status **Hardware does not match database** for this module and should show the action **Restore IBP**. Please perform this action for the new connection blade.
7. Once **Restore IBP** is successfully completed, plug in the LAN cables. Please make sure you connect them as they were connected to the old connection blade. For uplink sets with active ports and backup ports, you should first connect the active ports and then the backup ports.

When you have completed these actions, the new IBP module should work correctly.

15.2 VIOM address ranges

During installation of ServerView Virtual-IO Manager, you can select address ranges used for automatic assignment of virtual MAC and WWN addresses.

For virtualization of the MAC addresses of LAN I/O devices , you can choose from eight predefined MAC address ranges, which do not overlap (**MAC Address Range 1** to **MAC Address Range 8**). Each of these address ranges contains 8,000 MAC addresses. If such a range is insufficient, you can also select a range double the size using **MAC Address Range 1 and 2** to **MAC Address Range 7 and 8**. Each of these areas contains 16,000 MAC addresses.

These address ranges are defined as follows:

Address range	Start address	End address
MAC1	00:19:99:3E:D2:A1	00:19:99:3E:F1:E0
MAC2	00:19:99:3E:F1:E1	00:19:99:3F:11:20
MAC3	00:19:99:3F:11:21	00:19:99:3F:30:60
MAC4	00:19:99:3F:30:61	00:19:99:3F:4F:A0
MAC5	00:19:99:3F:4F:A1	00:19:99:3F:6E:E0
MAC6	00:19:99:3F:6E:E1	00:19:99:3F:8E:20
MAC7	00:19:99:3F:8E:21	00:19:99:3F:AD:60
MAC8	00:19:99:3F:AD:61	00:19:99:3F:CC:A0
MAC12	00:19:99:3E:D2:A1	00:19:99:3F:11:20
MAC34	00:19:99:3F:11:21	00:19:99:3F:4F:A0
MAC56	00:19:99:3F:4F:A1	00:19:99:3F:8E:20
MAC78	00:19:99:3F:8E:21	00:19:99:3F:CC:A0

For virtualization of the WWN addresses of Fibre Channel I/O devices, you can choose from eight predefined WWN address ranges, which do not overlap (**WWN Address Range 1** to **WWN Address Range 8**). Each individual address range contains 32,767,487 WWN addresses.

These address ranges are defined as follows:

Address range	Start address	End address
WWN1	50:01:99:93:ED:2A:10:00	50:01:99:93:EF:1E:0D:FF
WWN2	50:01:99:93:EF:1E:0E:00	50:01:99:93:F1:12:0B:FF
WWN3	50:01:99:93:F1:12:0C:00	50:01:99:93:F3:06:09:FF
WWN4	50:01:99:93:F3:06:0A:00	50:01:99:93:F4:FA:07:FF
WWN5	50:01:99:93:F4:FA:08:00	50:01:99:93:F6:EE:05:FF
WWN6	50:01:99:93:F6:EE:06:00	50:01:99:93:F8:E2:03:FF
WWN7	50:01:99:93:F8:E2:04:00	50:01:99:93:FA:D6:02:FF
WWN8	50:01:99:93:FA:D6:03:00	50:01:99:93:FC:C9:FF:FF

If you have an address range of your own that you wish to use for virtual MAC or virtual WWN addresses, then select it in the **Custom MAC Range** or **Custom WWN Range**.

You can also assign individual virtual addresses when defining a VIOM server profile.

Virtual addresses should always be taken from a reserved address range, otherwise you might find that you are using addresses from a vendor address range and this could result in duplicated addresses. The original addresses of a controller should never be used as virtual addresses in a VIOM server profile.

If you have several installations of the Virtual-IO Manager in your network, you must ensure that the address ranges used by two installations do not overlap. Otherwise addresses may be assigned several times to different systems, which results in duplicated addresses.

15.3 Creating diagnostic data

To activate the trace functionality of the ServerView Virtual-IO Manager service, please edit the file **ViomConfig.properties** in the directory **<ServerView Suite>\plugins\viom\Manager**.

On a Windows operating system this is typically the directory **C:\Program Files\Fujitsu\ServerView Suite\plugins\viom\Manager**.

Please modify the line **StartOptions=** by adding the option **--debug trace** (two minus signs!):

```
StartOptions=--debug trace
```

The trace file of the ServerView Virtual-IO Manager service is:

<ServerView Suite>\plugins\viom\Manager\logs\viom-manager.log

To also get full trace information from the module that configures connection blades such as IBP 10/6, IBP 30/12, SB11A or SB11, add the debug option **trace,driverdbg** to this line:

```
StartOptions=--debug trace,driverdbg
```

The ServerView Virtual-IO Manager service will then create a trace file for each connection blade configuration command. These files are written to the directory **<ServerView Suite>\plugins\viom\Manager\logs**.

The ServerView Virtual-IO Manager service creates a pre-defined maximum number of log files. Each log file can have a pre-defined maximum size. When the maximum number of log files is reached the oldest log file is automatically deleted. New log information is written to a new log file.

The maximum number of log files can be changed by editing the file **manlog4j.properties** in the directory:

<ServerView Suite>\plugins\viom\Manager.

In some cases it might be necessary to increase the number of log files in order to keep trace information for a longer period of time. If so, please modify the line **log4j.appender.DebugAppender.MaxBackupIndex=25** and set the property **MaxBackupIndex** to an appropriate value. Depending on the amount of requests sent to the Virtual-IO Manager service, it might be necessary to set this value to 50 or even higher if you want to keep log information for about one week.

The property **MaxFileSize** should not be increased, because it might make it difficult to load files into an editor.

After modifying the **StartOptions** property in the **Viom-Config.properties** file or the **MaxBackupIndex** property in the **man-log4j.properties** file , the service **Server-ViewVirtualIOManagerService** must be restarted. (Note: The display string of this service is: **ServerView Virtual IO Manager Service**.)

You can also activate the trace functionality for the VIOM provider of the ServerView Connector Service (SCS).

To do so you must modify the file <remote_connector_dir>**\ViomAPI.xml**.

Please modify the lines

```
<viom:debugging-level>0</viom:debugging-level>
```

```
<viom:tracing-level>0</viom:tracing-level>
```

by setting **debugging-level** and **tracing-level** to **127**. You will then need to restart the SCS service.

On a Windows operating system the default for the Remote Connector installation directory <remote_connector_dir> is:

C:\Program Files\Fujitsu\ServerView Suite\Remote Connector

On a Linux operating system the Remote Connector is installed in the directory:

/opt/fujitsu/ServerViewSuite/SCS

There you will also find the VIOM configuration file **ViomAPI.xml** containing, among other things, the name and location of the VIOM trace file:

```
<viom:logging-file>/var/log/fujitsu/ServerViewSuite/viom/viom-provider.log</viom:logging-file>
```

If you modify the VIOM configuration file, you must restart the Remote Connector for it to become effective.

On a Windows operating system the action **Collect Log Files** in the **Start** menu (**Start - All Programs - Fujitsu - ServerView Suite - Virtual-IO Manager - Collect Log Files**) creates a ZIP archive. This ZIP archive contains all log files from Virtual-IO Manager, VIOM database information, and log files from the ServerView ServerList service.

For how to collect diagnostic data on a Linux operating system, see the chapter "Collecting diagnostic information" on page 89.

To activate the trace functionality of the ServerView Virtual-IO Manager user interface, proceed as follows:

1. On the **Virtual-IO Manager** tab click the **Preferences** button.
2. In the **Preferences** dialog box, select the **Trace** tab. There select the option **Write trace messages to file** and enter a file name in the **File name** input field and change the **Max. file size** if necessary. The trace will be written to this file. If the maximum file size is reached, the trace file will be renamed and a new one will be used. The renamed trace files will have a number appended, up to a maximum of ten possible back-up trace files.
3. Reproduce the error.
4. Save the trace file(s) for later diagnosis.
5. Click the **Preferences** button again (see step 4) and deselect the option **Write trace messages to file**. If you do not do this, the trace will continue to be written; even after the user interface is restarted.

15.4 Event logging

The ServerView Virtual-IO Manager writes event logs if an error occurs (event type **Error**) or if the configuration is modified (event type **Information**).

Where the events are logged depends on the system on which the ServerView Virtual-IO Manager is installed:

- For Windows

  Virtual-IO Manager events are recorded in the **ServerView VIOM** event log of the Windows Event Viewer with source **ServerView Virtual-IO Manager**, **ServerView Virtual-IO Manager SOAP API**, **ServerView Virtual-IO Backup Service**, and **ServerView Virtual-IO License Manager**.

- For Linux

  Virtual-IO Manager events are output to the system log with source **VIOM-MAN**, **VIOM-LICENSE-MANAGER**, and **VIOM-BACKUP-SERVICE**.

The event logs for errors are self-explanatory and not listed here. Informational events that describe changes to the configuration are shown in the following table. The log entries contain more information (e.g. names of involved uplink sets or nodes) than listed in the **Meaning** column. If you are running Windows, this list of event IDs should enable you to filter the events for relevant entries.

Event ID	Meaning
50121	Session created.
50123	Session removed.
50221	Authentication for node set.
50222	Authentication for node changed.
50321	Node is now managed.
50323	Node is now unmanaged.
50421	Uplink set created.
50422	Uplink set modified.
50423	Uplink set deleted.
50521	Uplink added to uplink set.
50522	Uplink of uplink set changed.
50523	Uplink removed from uplink set.
50621	Network created.
50622	Network modified.
50623	Network deleted or network not native anymore.
50721	Profile created.

Event ID	Meaning
50722	Profile modified.
50723	Profile deleted.
50735	Profile assigned.
50736	Profile unassigned.
50821	IO-channel added to profile.
50822	IO-channel of profile modified.
50823	IO-channel removed from profile.
50921	Address range set.
50923	Address range unset.
51000	Power mode for node set.
51100	Virtual-IO Manager launched successfully.
51217	Node has been restored to ServerView Virtual-IO Manager database configuration.
51231	Configuration backup has been restored to node.
51321	Configuration of Virtual-IO Manager saved.
51323	Configuration file deleted.
51447	Authentication of specified user successful.
51521	License registered.
51523	License removed.

There are no event IDs available in Linux.